

The Indian Govt is Just Not that Into Freedom

Rs 100
ISSN 0974-1054

Free DVD
Live Linux Gamers

LINUX



THE COMPLETE MAGAZINE ON OPEN SOURCE

VOLUME: 07 ISSUE: 7 September 2009 116 PAGES ISSUE# 80

ForYou

Open Source in

Rocket Science

In 1997, two university students had a very simple plan for their graduate project—build a rocket capable of inserting a 1-kg pico-satellite into orbit.

[Read more on page 36.](#)

15 Games *on Linux Reviewed*

Who said they were non-existent? | 20

Sneak into Your Thumb Drive from
the Cloud | 88

Develop a Simple Download
Scheduler | 56

Building Appliances is Easy with
SUSE Studio | 26

Published by EFY—ISO 9001:2000 Certified

India INR 100
US \$ 12
Singapore S\$ 9.5
Malaysia MYR 19

**SUBSCRIBER COPY
NOT FOR RESALE**

Sale and purchase
of this copy is **illegal**.



The R7 rocket image on the magazine cover and here is a copyright of Sergei Korolyov. It's licensed under the Creative Commons Attribution ShareAlike 3.0 License.



FOR YOU & ME

- 18 Linux Gamers, Unite! It's Time to Frag!!
- 26 A Studio to Edit and Roll Out Your Appliance
- 30 The Perfect Desktop
- 33 KDE 4.3: It's the Desktop You Need, Almost
- 36 Who Says Linux Is Not (in) Rocket Science?
- 41 The GoI is Just Not That Into Freedom
- 48 Use a GPRS Handset as Your PC's Gateway to the Web
- 50 Video Editing on GNU/Linux. It's Easy With Kdenlive
- 53 Shenoy Systems: Addressing the Linux Talent Shortfall
- 64 GIMP for Beginners, Part 3: Toolbox Demystified



ADMIN

- 60 Tiered Storage Infrastructure: Storage, Server and Application Levels
- 70 Performance Tuning and Monitoring— Part 1
- 84 The Art of Guard—SELinux Logging, Part 5



CONTENTS

DEVELOPERS

- 88 Sneak Into Your Thumb Drive from the Cloud
- 91 Python for Research: An Initiation

GEEKS

- 44 Play With Your Partitions
- 56 Develop a Simple Download Scheduler
- 78 Save Bandwidth by Setting Up a Fedora Mirror



COLUMNS

- 96 The Joy of Programming: Typo Bugs
- 97 CodeSport
- 100 A Voyage to the Kernel—Part 16: Segment: 3.5, Day 15



REGULAR FEATURES

- 06 Editorial
- 08 You Said It...
- 10 FOSS Bytes
- 16 Q&A Section
- 94 Tips & Tricks
- 106 Linux Jobs
- 108 FOSS Yellow Pages



LFY DVD: Live Linux Gamers 0.9.5

A collection of 34 Linux games! It contains popular titles like Battle of Wesnoth, Extreme Tux Racer, Frets on Fire, Glest, Armagetron, Frozen Bubble and many more that shall keep you glued to your consoles.



LFY CD: KDE 4.3

+ *OpenOffice 3.1*, + *Firefox 3.5*, with *openSUSE 11.1* as base, and *codecs* fetched from *Packman*. A modern desktop with an office suite, multimedia codecs, wireless support and more. Well, you can have all that within two minutes!

Note: All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 Unported Licence a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence.

Editor

RAHUL CHOPRA

Editorial, Subscriptions & Advertising

DELHI (HQ)

D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 26810602, 26810603
Fax: 26817563
E-mail: info@efyindia.com

BANGALORE

No. 9, 17th Main,
1st Cross, HAL II Stage,
Indiranagar, Bangalore 560008
Ph: (080) 25260023; Fax: 25260394
E-mail: efyblr@efyindia.com

CHENNAI

M. Nackeran
DBS House, 31-A, Cathedral Garden Road
Near Palmgroove Hotel, Chennai 600034
Ph: 044-28275191; Mobile: 09962502404
E-mail: efychn2@efyindia.com

Customer Care

E-MAIL: support@efyindia.com

Back Issues

Kits 'n' Spares

D-88/5, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 32975879, 26371661-2
E-mail: kits@efyindia.com
Website: www.kitsnspares.com

Advertising

KOLKATA

D.C. Mehra
Ph: (033) 22294788
Telefax: 22650094
E-mail: efycal@efyindia.com
Mobile: 09432422932

MUMBAI

Flory D'Souza
Ph: (022) 24950047, 24928520; Fax: 24954278
E-mail: efymum@efyindia.com

PUNE

Zakir Shaikh
Mobile: 09372407753
E-mail: efypune@efyindia.com

HYDERABAD

P.S. Muralidharan
Ph: 09849962660
E-mail: efyhhd@efyindia.com

Exclusive News-stand Distributor (India)

INDIA BOOK HOUSE PVT LTD

Arch No. 30, below Mahalaxmi Bridge, Mahalaxmi,
Mumbai - 400034 Tel: 022- 40497401, 40497402,
40497474, 40497479, Fax: 40497434
E-mail: info@ibhworld.com

Printed, published and owned by Ramesh Chopra. Printed at Ratna Offset, C-101, DDA Shed, Okhla Industrial Area, Phase I, New Delhi 110020, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2009. All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under under Creative Commons Attribution-Share Alike 3.0 Unported License a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to publishing errors. Articles that cannot be used are returned to the authors if accompanied by a self-addressed and sufficiently stamped envelope. But no responsibility is taken for any loss or delay in returning the material. Disputes, if any, will be settled in a New Delhi court only.

Editorial

Dear readers,

I was not enthused when some of my editorial colleagues started branding the KDE 4.3 compilation, which features in the current *LFYCD*, as the ‘perfect desktop’. Well, for those who use GMail’s voice and video chat—it definitely did not seem like one. Could a desktop that does not facilitate the use of such popular tools, be called ‘perfect’? I wasn’t sure.

My colleagues, on the other hand, argued that the fault lies with Google for not making the video chat feature available in the Linux version of Firefox, thereby giving people another excuse not to switch to Linux. When you come to think about it, doesn’t Google always do this—release stuff for Windows first? But should we keep complaining or work towards getting that feature as a core functionality of an alternate application? Thus, when Google (or any other company, for that matter) finally makes the solution available to us, we can say, “Thanks for your hard work, but I’m already comfortable using something else.”

A lot of this has already happened in our open source world. Does anyone even care about Adobe Reader on Linux? Does anyone really depend on CyberLink’s PowerDVD player, any longer? There are many other examples, where the community did not wait for Windows-friendly brands to come up with their Linux versions. If Adobe stops releasing the latest Flash versions for Linux, the chances are, that by the time someone puts some sense back into those at Adobe, people would have switched to something else.

Anyway, so where is this leading to?

The good news—we don’t need to wait for Google to do us the favour, thanks to the recent announcement from the Pidgin project team. Their latest version 2.6.1 has “added XMPP Voice and Video support.” You will now be able to connect with GMail’s voice and video chat users from the Linux world too. According to Pidgin’s changelog, “Voice and video support with

Jingle (XEP-0166, 0167, 0176, & 0177), voice support with GTalk and voice and video support with the GMail Web client...” have been added in version 2.6.0—the following version (2.6.1) was released the same day with a few bug fixes.

Now, this is a real winner on the desktop Linux front. Another excuse not to switch to Linux has been deftly taken down by smart open source developers. Certainly a development that’s enabled the ‘perfect

Does anyone even care about Adobe Reader on Linux?

desktop’ to seem more ‘perfect’.

Somehow this reminds me of that Chumbawamba song that was pretty popular during the 1998 Football World Cup: “I get knocked down. But, I get up again. You’re never going to keep me down.” Don’t get me wrong. I don’t in any way imply that we’re playing catch up in all fields. In fact, we’re pretty much ahead on most fronts. Yet, let’s figure out our shortcomings, and then Konquer.

With that, I rest my case while you enjoy the ‘perfect desktop’, which is by no means perfect but will certainly make Vista users go green with envy.

Best Wishes!



Rahul Chopra
Editor, LFY
rahul@efyindia.com

You said it...



I've read your Fedora 11 review [www.linuxforu.com/reviews/fedora11-review]. I also subscribe to *LFY* and presently use Windows XP with Linux (Ubuntu 9.04 and Fedora 10) on my home PC. Though I have a medical background, I am interested in computers/IT, especially Linux and OSS.

About installation, I run into a problem every time I try to upgrade from Fedora 10 to Fedora 11 using the DVD (at the point when the installer examines the system for HD configuration).

Is there a way to upgrade from Fedora 10 to Fedora 11 using the DVD as a source (i.e., local repository)? I tried Googling for this, but almost all responses suggest upgrading using pre-upgrade through the Internet. However, I would prefer using the DVD, as it will save me considerable bandwidth and time.

—Jay Mistry, jaylinux@ymail.com

Atanu Datta replies: *Is there a reason why you don't want to go for a fresh install, instead of an upgrade from a previous version?*

It's always a much faster process if you install fresh, and then configure the RPM-Fusion software repository immediately to pull in the media codecs. This is because, as far as I remember, you'd need to uninstall all third-party software before using the DVD to upgrade F10 anyway. I won't be able to confirm this, though, as I'm not a Fedora user and only test it for features every six months or so after a new release.

If you are comfortable with IRC channels then #fedora is where you'll get more fine-tuned solutions from the regular users.



I am an engineering student and an avid distro hopper. Recently, I came across a distribution that really grabbed my attention. I am talking about Pardus Linux, a distro from Turkey. Now, this is not 'just another distribution'. Pardus Linux is somewhat unusual in that it is government-funded. It was created and is maintained by the National Research Institute of Electronics and Cryptology (UEKAE), an affiliate of the Scientific and Technological Research Council of Turkey. Pardus is being used in both the public and private sector in Turkey and the distribution is saving Turkish taxpayers millions of Euros.

Now, I think that this is a shining example of how the implementation of FOSS on a national level can be achieved. I am not saying that we don't have our own 'Indian' distros, but the thing is that here, the

movement lacks the drive that is required to achieve long-term goals. In my opinion, we can learn a lot from Pardus. This is precisely the reason why I'm requesting you to check out this distro, and if possible, bundle it with your next issue of *LFY*. A review on it would be highly appreciated.

—Rajit Singh, rajit.vikram.singh@hotmail.com

ED: *Thanks for sharing your valuable insight, Rajit. We completely agree with you. In fact, there's one person in the LFY team who regularly uses Pardus 2009 on his home desktop, besides lobbying for the distro to be installed on all the computers we use in the editorial department—anyway, most of the EFY group uses openSUSE and Fedora. The reason cited to our IT admin is "...support for most things, out-of-the-box."*

As for bundling the distro, we have already included the version 2008.2 in our May 2009 issue, besides a review in the magazine. We'll try and include it again soon, now that they have defaulted to the KDE4 desktop :-)

By the way, did you know that an Indian is also working with the Pardus team on a project very similar to SUSE Studio called Pardusman, as part of the Google Summer of Code programme? You can find more details at www.sarathlakshman.info/2009/08/06/progress-with-pardusman-on-web and we'll most probably carry an article on Pardusman in our October issue.



Any idea whether Android Market is accessible in India or not? I have a Magic handset.

—Mahendra, doshi_md@hotmail.com

Nimish Dubey replies: *Good to hear from you. Unfortunately, the Android Market cannot be accessed in India from the HTC Magic. HTC has been talking of a firmware upgrade that will fix this but we have yet to see it, which is rather sad because half the charm of using an Android handset is being able to access the Android Market.*

Please send your comments or suggestions to:

The Editor


LINUX FOR YOU Magazine

D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020

Phone: 011-26810601/02/03 **Fax:** 011-26817563

Email: lfyedit@efyindia.com, **Website:** www.openITis.com





Now comes Ubuntu Desktop Support Services, from Canonical

Canonical's Desktop Support Services include three offerings: Starter, Advanced and Professional. The Starter Desktop Service provides support for installation, setting up and basic functionalities such as Internet, creating documents and playing music and videos. The Advanced Desktop Service is for more experienced users who need help migrating files and settings from a previously used operating system, or need assistance with desktop



publishing and personal accounting. The Professional Desktop Service is for business users who have Ubuntu as their main environment. Installation support ensures that the Ubuntu machine is set up on the corporate network and integrated into existing IT services.

Red Hat extends its partner ecosystem

Red Hat has launched the enhanced Partner Program initiative. The Red Hat Partner Program will now include a Premier Business Partner level, in addition to Advanced and Ready classifications. The program will also feature infrastructure, middleware and virtualisation specialisations. Additionally, the redesigned program is intended to enable Red Hat's partner community to collaborate and innovate to create complete solutions that solve customers' IT challenges.

Google Wave is now open source

Internet search giant Google has open sourced two components of its new-age communications platform, Google Wave. Launched in May this year, Google Wave is a Web platform that combines e-mail, instant messaging, social networking and document collaboration.

The two components that have been open sourced are Operational Transform (OT) code along with the underlying wave model, and the client/server application prototype to drive the interest of third-party developers. The company is promoting developer interest for implementations and hence making it easier for users and businesses to customise applications.

All of this —nearly 40,000 lines of Java code—is available under the Apache 2.0 licence. Google has also announced that the protocol specification, white papers and API documentation are being made available under a Creative Commons attribution licence.



tr.im URL shortening service opens up

The tr.im URL shortening service, currently run by Nambu, is to become a community-based service and the source code to the URL shortener is to be released under an MIT licence. The plan is that by September 15th, Eric Woodward, president, Nambu, will donate the tr.im domain to the community, release the source code as open source and offer link-map data to all.

Woodward will also personally guarantee shortfalls in the new tr.im's operating expenses, 'indefinitely'. tr.im will then accept donations to help meet those expenses and, to aid transparency, all statistics on usage, operating cash flow, redirects and creation counts, will be published.

VMWare buys SpringSource for \$420m


VMWare has acquired SpringSource for \$420 million in a deal that includes the developers of the Spring framework for Java applications, application server technology, the Groovy and Grails experts from G2One and the Hyperic management tools that SpringSource acquired in May. VMWare intends to use the SpringSource technology to create a fully integrated 'Platform as a Service' (PaaS) solution, combining VMWare virtualisation and management with a complete Java stack, from servers to frameworks and languages, with the Hyperic management tools integrated into the platform as a first class citizen.

Mono releases beta of Moonlight 2.0

Following a series of preview versions, the Mono Project has now released a beta of Moonlight version 2.0, the free Silverlight implementation. The software, which is available to download as a Firefox plug-in, allows WMA files and multimedia content implemented using Microsoft's Silverlight, to be played under Linux. Streaming is said to be substantially improved under Moonlight 2.0.

This version is mostly licensed under the GNU LGPL 2, except for the Microsoft Silverlight Controls and Unit Test Framework, which is under the MS-PL licence, and some managed code and tools that are under an MIT licence.





Arch Linux 2009.08 introduces kernel 2.6.30

Arch Linux release engineer, Dieter Plaetinck, has announced the release of the project's official 2009.08 installation images. The new version features the 2.6.30.4 Linux kernel (which brings in support for ext4 filesystem) and version 3.3.0 of its Pacman package manager.



In addition to a wide variety of bug fixes, boot device detection has been enhanced and Aufs2 (advanced multi-layered unification filesystem version 2) replaces the

UnionFS used in previous versions. The AIF (Arch Linux Installation Framework) installer also includes several updates, adding support for automated installations using config files, better date/time setting with NTP, various usability improvements and a new framework for users to build their own installation procedures.

ChromeDevTools for Eclipse users

Google software engineer Alexander Pavlov has announced the release of ChromeDevTools for Eclipse users. ChromeDevTools enables JavaScript debugging inside the browser over the TCP/IP-based ChromeDevToolsProtocol using Eclipse and consists of two components: an SDK and a debugger.

The SDK includes a Java API to enable communication over TCP/IP with Google Chrome or its Chromium counterpart. The debugger, a plug-in for Eclipse that uses the SDK, lets developers debug JavaScript running within Chrome from the Eclipse IDE. Using ChromeDevTools, developers can inspect variables, set breakpoints and evaluate expressions. ChromeDevTools requires version 3.4.2 or later of Eclipse and Google Chrome or Chromium version 3.0.189.0 or later. ChromeDevTools is released under a New BSD Licence.

Perl 6 due in the spring of 2010

In a recent entry on his blog, Patrick Michaud has again confirmed the news that has been circulating for a while — that Perl 6 will be released in the spring of 2010 (with all of the signs pointing to April). Michaud is the 'Pumpking' (keeper of the pumpkin patches—or release manager) for the Rakudo Perl 6 compiler, designed to run on the Parrot virtual machine.

The new spring version of the Perl scripting language will not include all of its functions, but it will be useable. According to Michaud, the release will be called 'Rakudo *' (Rakudo Star). Michaud says that language development is more of an 'evolutionary process' resulting from the interplay of language development and implementation efforts.

Rakudo Star will be released as version 1.0 in spring 2010, and the developers hope that the 1.0 tag will signify to users that, while it may not be complete, it is a usable implementation of Perl 6.

OpenGL 3.2 specification released

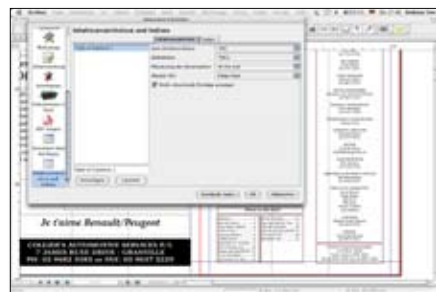
The Khronos Group has announced OpenGL 3.2, the third major update in 12 months to the most widely adopted 2D and 3D graphics API for personal computers and workstations. This new release continues the rapid evolution of the OpenGL standard to enable graphics developers to access cutting-edge GPU functionality across diverse operating systems and platforms. The full specification is available for immediate download at www.opengl.org/registry.

OpenGL 3.2 adds features for enhanced performance, increased visual quality, accelerated geometry processing and easier portability of Direct3D applications. In addition, OpenGL and other standards within Khronos, including OpenCL for parallel compute, OpenGL ES for mobile 3D graphics and the new WebGL standard for 3D on the Web, are being coordinated to create a powerful graphics and compute ecosystem that spans many applications markets and devices. The installed base of OpenGL 3.2 compatible GPUs already exceeds 150 million units.

Scribus 1.3.5 introduced

Following many testing cycles in the past two years, the Scribus Team has finally announced the release of Scribus 1.3.5. The team considers this version to be fairly stable and ready for many real-world use cases. As proof, the cover of the Official Scribus Manual has been created with a release candidate of 1.3.5. While there are still a few chinks and bugs in the latest version, the team will continue to implement new features and improvements in a new experimental branch. Once tested and confirmed as working flawlessly, they will be backported and become part of the 1.4 milestone.

Major changes in the latest release would be: a native OS X version (available as a .dmg file), based on Qt 4, drag and drop improvements, optional embedding of PDF and EPS files, a new line editor for custom line styles and a lot more. Furthermore, about 1,100 feature requests and bugs, both large and small, have been resolved.





KNOW HOW



❗ I have heard that all computers with an on-board LAN card have a unique MAC address. Please let me know what this MAC address is and how to check it on my desktop.
—Nilima Sharma, Agra

MAC is an acronym for Media Access Control, which is a hardware address that uniquely identifies network adapters or network interface cards (NIC). This is assigned by the manufacturer for identification, and used in the Media Access Control protocol sub-layer. It is also known as an Ethernet Hardware Address or physical address. To know the MAC address of your NIC, in Linux you can check the output of the `ifconfig` command. It displays information of all network interfaces along with their MAC addresses. Run the following command as the root:

```
ifconfig | grep HWaddr
```

❗ I have a HP Pavilion DV2401TU laptop with a

Broadcom 802.11b/g bcm4311 wireless adaptor. I tried to configure it on Linux (Mandriva, Red Hat, SUSE, Linux Mint) using `ndiswrapper` and patches from the Net, as per instructions. But nothing has worked till now. Please help me out here. What about the availability of wireless adaptors that are fully compatible with Linux? Kindly provide an apt solution for the problem.
—vik01@itimes.com

The Broadcom wireless LAN card is not supported out-of-the-box in Linux but you can surely configure it. Connect to the Internet using a wired interface and install some software.

For OpenSUSE, a 1-click install YMP file (YaST Metapackage file) is available from Packman. To install this, go to <http://packman.links2linux.org/> and search for `broadcom-wl`. Click on the '1-click install' icon and follow the on-screen instructions. To get it working on Mint 5, you can follow the instructions given at www.linuxmint.com/wiki/index.php/Broadcom_bcm43xx. As for Fedora and Mandriva, it's available in the RPM Fusion and Mandriva non-free repositories, respectively. You can find a list of Linux-supported wireless adaptors at <http://linux-wless.passys.nl>.

❗ I have an Intel Core 2 Duo system with 2 GB of RAM and a 160 GB HDD. I had installed Mandriva along with Windows XP. The system was working fine till I tried to install Fedora 11 from a live CD. After installing it, I am unable to boot any of my OS. I think there is some

problem with my MBR, as now whenever I boot, I get a message that says "Non-system disk or disk error. Replace and strike any key when ready". I have checked my BIOS and it recognises my HDD. Is the error due to HDD failure or is it due to some problem with the MBR?
—Milind Soneja

It seems that you have a bad boot loader. Just try installing GRUB from any live CD.

❗ I have installed Fedora and have created different users and groups. Is it possible to add a new user to a specific group? Also, let me know how to change an existing user's group.
—Vikas Jha

To add a user to a specific group, you need to use the command given below:

```
# useradd -G <specific_group_name> <username>
```

For example, if you have a group called `accounts` and you want to add a user account to it, then, issue the following command:

```
# useradd -G accounts accounts
```

This will add the user `accounts` to the group `accounts`, provided there already exists a group called "accounts". Otherwise you need to create the group first.

To modify an existing user you can use the following command:

```
# usermod -a -G <group name> <user to be modified>
```



Linux Gamers, Unite!

It's Time to Frag!!

Welcome live.linux-games.net. They bring you a collection of 34 games, safely tucked into a shell of Arch Linux. Be it puzzles, arcades, shooters or strategy, the collection has it all. And if you are a true gamer, you would not get distracted by this article...till you'd gutted it out on the included DVD!



Games on LLGN

- Armagetron Advanced
- AstroMenace
- Blobby Volley
- Chromium BSU
- Extreme Tux Racer
- Foo Billard
- Frozen Bubble
- HedgeWars
- LBreakout 2
- Pingus
- Quadra
- Secret Maryo Chronicles
- Tee Worlds
- World of Goo (demo)
- X-Moto
- BTanks
- Frets on Fire
- Glest
- Mania Drive
- NeverBall
- NeverPutt
- Nexuiz
- OpenLeiroX
- Sauerbraten
- Scorched 3D
- Super Tux Kart
- Torcs
- Tremulous
- UrbanTerror
- Warsaw
- Warzone 2100
- Battle of Wesnoth
- World of Padman
- Widelands

“Games on Linux? What on earth would games ever do, on a Linux box? All the Linux users are anyway hardcore geeks who do not give a fig about games!”

Yes, those were the exact statements I got from one of my Facebook contacts, when I put up a status update mentioning that I was reviewing a few Linux games for this month's issue of *LFY*. Belonging to the guild of professional gamers myself (I play only DotA these days, professionally), I always had to keep Windows installed in my box, be it on dual-boot or in a virtual box. Now, I can finally see a bright light at the end of what I had considered to be a tunnel with a dead end.

linux-gamers have definitely scored the first blood in this arena. The very idea of bundling a set of awesome games in almost all genres, mind you, in a DVD for everyone to play (or try out, in the case of those with Intel graphic chipsets—shall come back to that later) and cherish all their lives...

Well, live.linux-games.net (LLGN) is a project based on Arch Linux that lets you simply pop a DVD or CD into your x86 computer and start running games. That's it.

No configuration. No installation of games. No desktop distribution problems to wrestle with. You simply insert the DVD, boot up your computer and start playing. No fuss, no muss. (Check out the table along side for a whole list of the games.)

Hardware requirements

While LLGN has both the CD and DVD formats, the CD is targeted at leisure gamers and those who have a very basic machine. The minimum requirements suggested for the DVD we have bundled with this month's issue, are:

- i686 capable architecture (yes, that means you can run it on Intel's Macs too)
- 512 MB RAM
- Video card with 3-D acceleration

However, something needs a special mention here. LLGN's official website states that they haven't 'extensively' tested the live DVD on ATI and Intel graphics cards. My personal recommendation is, if you are a hardcore gamer and do not have a NVIDIA card, stay away from the live DVD, read the games' reviews in this article, and just download the ones you fancy, off the Internet.

So let's dig into the games!

Battle of Wesnoth



Genre: RTS (real time strategy)

Game play

This is one of the less common strategy games that concentrate heavily on combat skills, with hardly any emphasis on resources. While the 'turn-based' game play might be a bit of a turn-off, it also boasts of a unique set of mechanics that allow for a broader range of conflicts. Moreover, there is no tech tree, or even player-constructed buildings. The only resource is gold that can be collected by occupying more villages.

The terrain plays a key role in game play. In most cases it is the sole factor in how likely a unit's attack will actually land on its target. Also, different terrains offer differing amounts of protection for different units. Mountain tiles offer great protection for Elven fighters, but almost none for mounted units like knights. Sometimes the player has to make choices between chasing down decimated units or heading for the high ground in preparation for the next attack wave. Often, unit positioning is such a key element that the game feels more like chess than anything else.



The battles are often close and exciting, despite the turn-based nature of play. Sometimes entire conflicts can rest on the fate of a single unit. Sometimes smaller armies can defeat significantly larger ones with good tactics. The best thing about this game is that cheap tactics are non-existent or ineffective. There is no opportunity for tower raping or siege tank abuse. Even in the most dire situations, hope is not lost if one's army can escape to find cover and regroup.

Video

The graphics in this game are 2D sprites, with nothing but sliding around and fancy animation to serve as eye candy. That being said, they are some of the most polished graphics of any tile-based 2D game. The units and terrain are beautifully drawn in a cohesive style that lays emphasis on being cute over trying to pretend there is a third dimension involved. Tile transitions are seamless, thanks to what I believe is a careful blend of hand crafted borders and use of alpha transparency. The magic effects are brilliant. The combat animations are simple, but clever. Overall, there is little that could be done to make the graphics of this game any better than they already are.

Audio

The audio is very engaging and actually manages to engulf you into the game play if you have even a 2.1 speaker set. The sounds of combat are very realistic, my favourite being the clanking of the swords between the Elven archers and the Orcs.

Download URL: <http://www.wesnoth.org/wiki/Download>

HedgeWars



Genre: Action

Game play

Hedgewars is a multi-platform, free, turn-based, blast-'em-up in the style of Worms. It features a wide range of weapons, multiple maps, and excellent graphics and audio. It offers single player, multi-player, and (in theory) 'multi-player over the Internet' modes. It is a very well made multi-platform Worms clone that is free, fun, and offers some good production qualities.

Download URL: <http://linux.softpedia.com/get/GAMES-ENTERTAINMENT/Arcade/Hedgewars-28062.shtml>





Frets on Fire

Genre: Musical

Game play

Frets on Fire is an open source clone of the ever-popular PlayStation 2 music/rhythm game Guitar Hero. It's a free download and is available for Windows, Linux and the Mac, though the Mac version is experimental.

Just like Guitar Hero, Frets on Fire features a handful of songs



and allows you to play the guitar sections with just a few keys. Instead of using a separate controller to mimic the guitar, Frets on Fire lets you pick up the keyboard and jam right in front of your desk. Best of all, a built-in editor lets you tab any song, spawning a growing community of music-modders on the Web.

Rhythm games add an entirely new dimension of fun to gaming in the form of aural rewards. Frets on Fire takes that to a new level by allowing you to re-create familiar tunes with nothing more than your keyboard. The reward isn't getting a high score, it's listening to the song playing on your speakers and knowing you were responsible for that. The feeling you get when you effortlessly land an impossible

combination is amazing.

Also worth mentioning is the extremely entertaining tutorial. Your host, Jurgen Guntherswarchzhaffenstrassen, walks you through the game and proves that he is, indeed, much cooler than you are.

Video

The game doesn't really need anything flashier than the cartoon-isque animation it already has.

Audio

This is the best part. It's you who gets to decide how the game sounds, mistah! Mine sucked big time, though.

Download URL: <http://fretsonfire.sourceforge.net/>

Nexuiz



Genre: MMORPG
(massively multi-player
online role playing game)

Game play

If you have played Doom, Quake, Unreal Tournament, or any of their various free mods/clones, then you already have a good idea of what to expect from Nexuiz. Like all those games before it, Nexuiz is simple, dumb fun. Turn the game on and within 20 seconds you will be fragging people hundreds of miles away.


The community surrounding the game is not enormous, but the 500,000 plus downloads ensure that the servers will always be well stocked when you are ready to play. A small collection of North American clans has started to emerge, and the map-makers always manage to keep a steady

collection of maps flowing into the community.

However, the choice of guns is a lot less and the maps took ages to download when we tried playing it online. What's great was that there were a lot of servers to play in.

Video

When it comes to visuals, Nexuiz seems to have something for everyone. The lowest settings will easily run on very modest computers, while the highest settings make for a very nice light show. These niceties do not compare to the visuals of their modern commercial counterparts, although most are willing to forgive this since the game is entirely free and is based on a modified version of the Quake 1 engine (which came out in the mid-90s).



Audio

The sound effects really help immerse the player in the action, and can be a great help during intense 1v1 battles. However, the music is weak, derivative techno that you will surely switch off in favour of your own tunes, or just the sound effects.

Download URL: <http://www.alientrap.org/nexuiz/>



Blobby Volley

Genre: Arcade

Game play

Blobby Volley is an action-filled, multi-player, addictive volleyball game for 1 or 2 players. The players can go head-to-head with each other either on the same machine, on LAN or even on the Internet. You can also play against the bot or get two bots against each other if you wish to be a lounge lizard and just enjoy the action. Personally,



I think you do not have to be a volleyball fan to be able to play or enjoy this game.

Video

Just a bunch of 'jelly babies' created using vector graphics. No rocket science there.

Audio

Almost non-existent.

Download URL: <http://linux.softpedia.com/get/GAMES-ENTERTAINMENT/Arcade/Blobby-Volley-2-20021.shtml>



Armagetron

Genre: Arcade

Game play

The idea is simple. You have to keep building walls and outrun your opponent by not bumping into either your wall or your opponents.

While the game has options of playing against the bot, the game truly shines when it is played against a number of human players (LAN or online).

Video

In one word—average. But then, it is the game play that will keep you immersed in it for hours on end. The quality of video will be the last thing on your mind.

Audio

While the general sound effects are nothing exemplary, bumping into the walls does sound very realistic.



Download URL: <http://www.armagetronad.net/downloads.php>



Chromium BSU

Genre: Arcade

Game play

While this is just another classic overhead-scrolling space shooter, it is a game with a twist. You cannot let any of the alien ships past you; each one will cost you a life. Moreover, the ammunition is also limited. This is what makes the game very challenging, almost impossible.

Video

The animation of the power-ups and new ammunitions is very

intriguing. Other than that, it's as good as AstroMenace.

Audio

Good enough to keep you engrossed into the game.

Download URL: <http://linux.softpedia.com/get/GAMES-ENTERTAINMENT/Arcade/Chromium-B-S-U-9998.shtml>





Glest

Genre: RTS (real time strategy)

Game play

Glest is a 3D real-time strategy game (or RTS) set in pseudo medieval Europe. The game, as of version 2, features two factions, Tech and Magic. Each faction also has its own unit types, buildings and tech tree (i.e., more stuff to do later). As is typical with RTS games, there are resources to gather and mine.

The interface is quite clutter-free with just the essentials on the screen. At the top of the screen is a bar showing the resources available, like gold, stone and wood; and on the right are your faction management options. It's here that you get to build, upgrade and research new stuff. Rolling over each option tells you what it is and what resources are required.

However, we found a few glitches in the game. The tutorial

was a bit inaccurate. While we were asked to build a certain structure, we found we did not have enough resources and we weren't yet taught how to mine for or gather resources, by that time. Moreover, the game is extremely slow and the lighting is also pretty low, at any time of the day or night. Notification of buildings built or a worker being idle would have been very useful.

Video

Glest's theme is mainly a forest. Each unit has its own animations rendered in 3D. The ogre giant and battle walkers dwarf the other units.

Audio

Sound is probably one of the most underused elements in many open source projects but Glest's effects really boost the game's appeal. The music



is well composed and quite pleasant too. Each unit has its own response when you click on it or send it to do a task. The stone golem and magic armour are particularly good with grave, hollow voices.

Download URL: http://sourceforge.net/projects/glest/files/glest%20source/3.2.2/glest_source_3.2.2.zip/download



Frozen Bubble

Genre: Arcade

Game play

Frozen Bubble is a pop-cap style game that puts the user into a world of coloured balls, ice and penguins. While that might sound a bit silly, that's the whole point. The game works on the same principle as several very popular pop-cap style games such as Dinosaur Eggs and even the classic, Puzzle Bobble that started it all. The concept of Frozen Bubble is simple. The screen is filled with coloured balls that hang down from a gradually descending press plate. Your objective is to aim your igloo cannon at these balls and fire your own coloured balls

at them. If you match three in a row of the same colour, any attached coloured balls of the same colour, or any below them, immediately fall. It's almost like playing Connect Three in reverse. Of course, if the balls that haven't fallen yet touch the igloo, the game's over.

While the game might seem custom made for kids, it will



work as a great stress buster if you manage to sneak the live DVD into your office! Moreover, you can also play against other players on the Internet, and show off your score to the community.

While we could easily get connected to the online server, there were only seven games and three players online. Moreover, what seemed a bit weird was that, other than me, everyone was named FINALFANTA.

Download URL: <http://www.frozen-bubble.org/downloads/>

Words of Caution: EXTREMELY ADDICTIVE!!



LBreakout 2

Genre: Arcade

Game play

Lbreakout2 is a fun arcade game, where you use a paddle to deflect balls, aiming at bricks in the air. Like many other games of its type, Lbreakout features many power-ups and downs, brick types, and other great bonuses.

What we loved the best are the extensive balloon tips. If you wish to get it in full screen, you need to press the 'F' button. The game has a



total of 31 levels and can be played online too. However, we could not

get connected because the server's 8000 port seemed to be blocked.

Video and audio

The graphics is strictly decent. Everything is textured and appropriately detailed. The audio too is good, including the bouncing and hitting sound effects.

Download URL: <http://lgames.sourceforge.net/index.php?project=LBreakout2>

Secret Maryo Chronicles

Genre: Adventure

Game play

Secret Maryo Chronicles is an open source PC game designed around the original Super Mario World side scroller that takes what was great about that old console game and brings it to the PC, dusts it off and tries to improve on it. While still not completely polished or fully matured, it does offer you a lot of gaming fun. So let's look at the little that it does offer you.

In the game, you play as a young kid known as Maryo. One of the things you will quickly notice is the strikingly familiar backgrounds, characters and objects in the game. As I said, the game emulates Super Mario World quite well. The thing I

didn't like was how the game played out. It feels a bit loose and unpolished, as I mentioned earlier. However, the excellent graphics and eye candy at least partially make up for that. It's not like the game isn't playable, because it fully is. It's just that it plays in such a different way that you have to tinker a bit to get a full grasp of how things work before diving headlong into the game.

The levels on Secret Maryo Chronicles though are one of the best features of the game. I found them to be both challenging, and fun to explore. Unlike Super Tux and most of the original Mario games, you can actually go backwards on the map, all you want. You're not locked into one point on the map once you pass a certain spot. You can go all the way back to the beginning, run around, explore, or do what you like. While not

necessary, it does lend itself to being rather fun.

On the down side, the loading time of the game was very high and the controls were absurdly fast. As for the game play, we discovered we needed to jump on a tortoise twice to make it slide. If we touched it, we were dead.



Video

The quality of the graphics is actually good. While the old Super Mario World was in 16-bit Nintendo splendour, Secret Maryo Chronicles is in simulated 3D. I

say simulated because while it's not actually 3D, it looks like it is.

Audio

While the music was pretty good, it did not really sync up well with the in-game actions.

Download URL: http://www.secretmaryo.org/index.php?page=game_downloads



Extreme Tux Racer



Genre: Racing / Adventure

Game play

I guess, after SuperTux, this would be the most popular game on the Linux platform. Tux Racer is simple and very entertaining to play. Here, you control Tux as it slides on the mountain, to catch as many herrings as possible. Also, you will have to finish the race in time, or you lose. Moreover, the game needs you to jump over glaciers, glide in the waters and also perform 'tricks'.

As with any coin, this one has another side too.

Changing video resolutions screwed up the alignments. The game was also a bit buggy since it got me stuck into the mountains thrice in a row (probably punishing me for racing Tux over everything I possibly could). Moreover, the model of Tux could be made a bit more realistic. It seemed very 'plasticity', in the game.



Download URL: <http://www.extremetuxracer.com/?download>

Pingus



Genre: Puzzle

Disclaimer: This is undoubtedly the best game I have played, in the whole DVD. So, expect my review to be a bit fanboy-ish.

Game play


It has been scientifically proven that breaks at regular intervals increase productivity and there could be no better game (other than Desktop Towers Defence) than Pingus to fit into that bracket. What's more, Pingus is so addictive that it definitely manages to keep you rooted to your seat, scenario after scenario.

This game is basically an adaptation from the hugely

popular Lemmings game, i.e., combining various skill sets to attain a certain goal. Basically, you have a specific number of penguins that have to be saved to accomplish certain missions. Penguins start showing up at a certain point of the map, and you have to carefully guide them until they reach the exit.

Every mission has certain abilities that you can equip the penguins with, in order to pass obstacles that stand in their way. For instance, you may sometimes need to equip your penguins with floaters (little propellers mounted on their head), so that when they reach a gap, they gently land on the ground, and not hit it at great speed and die. You can also make the penguins jump over gaps to reach a specific point, or better, build a bridge over the gap and never have to worry about making sure every penguin jumps over.

The only part about the game that I hated was that it had too few levels (just 8 Halloween maps).



Video

The graphics of Pingus is satisfying but, given that it's a 2D game designed more for thinking than to please the eye, it can be deemed quite good.

The entertaining aspect comes not from superb landscapes or appealing effects, but from the contentment you feel when you finish a mission. So much so, the animation rendered to the penguins when they are enabled with certain abilities are also quite cute. I will always vote for the 'basher' animation.

Download URL: <http://pingus.seul.org/download.html>



AstroMenace

Genre: Arcade

Game play

AstroMenace is a brilliant 3D scroll-shooter that asks you to guide a spaceship, and destroy enemy ships and alien invaders in a fierce space battle. What makes it most interesting is the option of upgrading spaceships, advanced weaponry, missiles and armour. It also comes with pilot profiles so that your younger brother doesn't start fragging the aliens and skew your streak! And yes, AstroMenace definitely reminds us of Tyrian, a lot.

Video

The 3D graphics are a treat for any seasoned arcade player. The various kinds of ammunition and



spaceships have been designed very well and the brilliant rendering makes the gaming experience a great one.

Audio

Nothing ground-breaking; but neither is it too boring, as is normally the case with such scroll-shooter games.

Download URL: <http://www.viewizard.com/download.php>



Just a few more words...

So, as is evident, LLGN has done a great job in compiling the games for us to review. I say 'review' because, being a live DVD, you will not be able to save the game data for the future. What that means is, if today you finish the first 7 levels of a game and then plan to continue it tomorrow, the very moment you reboot, all that data gets lost and you end up playing the same levels all over again. Moreover, there will be a lot of games that simply refuse to run on Intel graphics. So serious was the issue that we actually had to find a specific NVIDIA machine to test drive all the games.

The DVD comes with a very basic installation of Arch Linux, with IceWM as the desktop environment. It comes preloaded with Gran Paradiso (a nightly build of Mozilla Firefox 3.0.11), the terminal and a network manager. While the distro detected our wireless cards seamlessly, if it doesn't do the same for you, a LAN card should come to your rescue.

Summing it up, while the collection of games is really addictive and fun, open source games are still aeons away from the professional games that get developed for the Windows and Macintosh platforms. Personally, I feel it is the community's fetish for open source games that doesn't lure the game publishers to give the Linux platform a serious look. While it might be more of an economic aspect, LLGN helps us show the world that Linux, as a platform, is as suited for gaming as anything else, if not better! **END** 

Mania Drive

Genre: Arcade

Game play

This car game is extremely addictive and frustrating at the same time. The idea of the game is to challenge the player against a variety of tracks that have really sharp turns, jumps and 360-degree hoops. On the whole, the game play is very smooth and reminds me a lot of one of the most popular games in this genre, Track Mania.

However, the tracks are few and once they are done, the only fun left would be to directly

pit yourself against human opponents, online. What makes it more challenging is that you've got to finish every single track within 'gold' time, else you lose. No mercy at all!

Video

While the background is too bland and so is the car, they do not dampen the spirit of the game even a single bit.

Audio

Mania Drive features eight songs that are all licensed under the Creative Commons and can therefore be freely distributed within the game. Most of the ogg-files can be categorised as rock music and that perfectly fits the fast, dynamic game play.

Download URL: <http://maniadrive.raydium.org/index.php?downloads=yes>



By: Sayantan Pal

An avid Twitter user and a social media enthusiast, the author is a passionate blogger and a professional gamer too. He also feels compelled to be opinionated about anything that comes his way, be it Linux distributions, our marketing strategies, table etiquettes or even the fabled Ramsay movies!



to Edit and Roll Out Your Appliance

The Web-based SUSE Studio service makes the job of rolling out your own GNU/Linux-based distro or appliance child's play.

Typically, for a systems integrator (SI) or even a systems administrator, installing, configuring and then setting up an operating system as per the requirements is a time-consuming job, to say the least. It involves adding all the extra essential apps, deleting all that's not required, configuring other apps, etc.

What if you're an ISV? You offer customers an application but don't have an OS. You sell (support) or give away your application and whether it works or not depends on the person installing it on top of an OS on the other side of the fence. Still, you'll have to make sure that it runs properly on that OS, while sitting on this side. Or you could dictate the terms under which your application has to be used—that the customer needs to run a particular OS, for instance—which, most often, is the case anyway. The result? Some of your potential customers couldn't care less about your product.

Or, maybe you're just a regular Linux user who needs to create an OS that gives you all the applications and settings out of the box, just the way you like it. Because, configuring a default installation *your* way takes too much time. And if you need to do that every few months after a new release, it becomes quite tiresome.

What if all of you had a way to simply customise an OS with only the applications you needed and nothing extra? If, as the producer of a third party software (an ISV), you could just add your application atop the base OS, build an ISO or a virtual appliance, test it, then download and use/distribute it?

I'm sure a lot of you have probably been doing precisely this for some time—but the process isn't as straightforward as you'd want it to be. A rather wizardish approach would make things easier.

Well, I don't know what the team behind a Web service called SUSE Studio [susestudio.com] had in mind—maybe they had the same plans—but it seems this

service just ended up solving a lot of issues, including the ones mentioned above. All at one go!

So, now that we've understood the target audience for this sort of service, let's get started, shall we?

On your marks

Well, to begin with, it's a Web-based service! The good thing is you can use your OpenID (or even Gmail and Yahoo! IDs) to sign in. It doesn't force you to sign up for a Novell account. The bad news is that the service is by invitation only, for now. Remember when Gmail was launched in 2004? Send them a request, and keep your fingers crossed—although the report is that most get to create an account within a week.

You'll start off with a base template. The service provides three distros as options for a base OS—openSUSE 11.1, SUSE Linux Enterprise (SLE) 10 and SLE 11. Each of these distros have six templates to choose from:

1. JeOS (Just-enough OS) – Well, as the name suggests, it's just enough to be called an OS, suitable for minimalistic appliances.
2. Server – A text-only base
3. Minimal X – X Server + IceWM
4. GNOME desktop – base system + GNOME
5. KDE4 desktop – base system + KDE4
6. KDE3 desktop – base system + KDE3

After this, you'll have to make a decision between 32- and 64-bit architectures.

I'll take the example of the KDE4 distro we're bundling with LFY this month, to run you through the process. I selected openSUSE 11.1, KDE4 desktop and the 32-bit architecture. Now give a name to your appliance—mine's KDE 4.3 Live CD.

It's time to put on your editor's hat, because you now enter the wizard to start customising your distro/appliance. Click on the **Software** tab on the Web page. This is where you'll do most of your work. This section is divided into three: **Software sources**, **Selected software** and **Search for software**. See Figure 1 for the complete picture.

Under **Sources**, only openSUSE 11.1 and Updates repos are available, to start off with. The selected software shows the software that's included, by default, based on the template you've started off with, and the **Search** section is where you're to locate what all is available in the repos. It is all interactive. One good thing is that when you enter a search term (by package name or pattern) it lists what's available, ranked by popularity. It lists the name of the application, its version, repository in which it is available, installed (not the download) size, and an action button to add (or remove already selected) applications.

Your first job is to identify what exactly you want. My job was to create 'The Perfect Desktop'. For that I needed the recently-released KDE 4.3, the latest version of OpenOffice.org and Firefox, besides support for multimedia codes and certain forms of drivers, out-of-the-box.

The default repositories offer older versions of all the

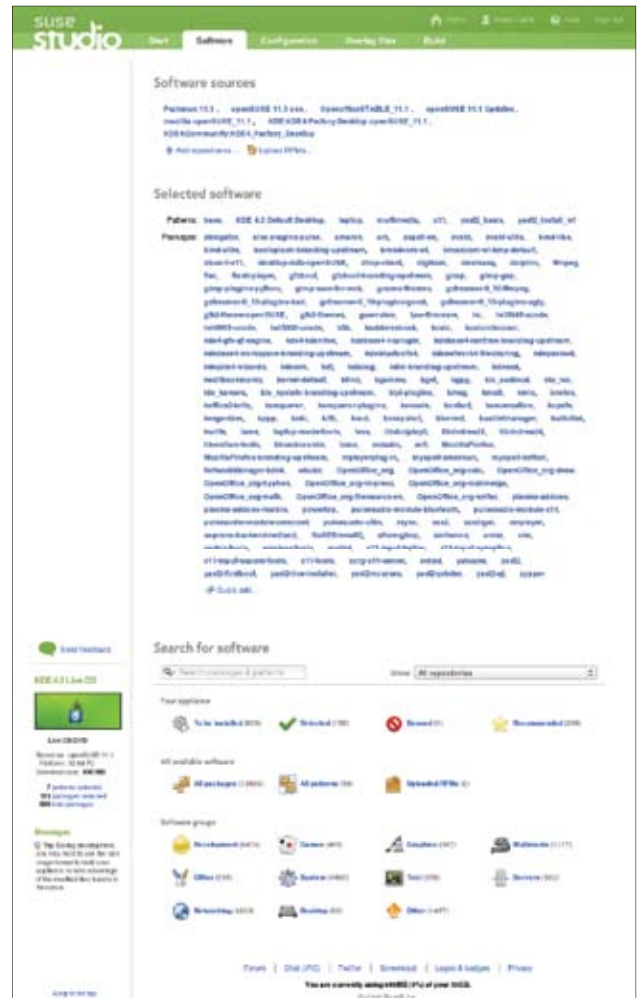


Figure 1: Under **Software** section we add applications we require

essential applications, and the codes are plain missing in them.

Go to the **Software sources** section and click on 'Add repositories' (Figure 2). Here, you can either search for a repository if you know its name or search for a package and see which repository makes more sense to add, based on the popularity ratings and number of packages available. For example, a search for KDE 4.3 pops out two results, out of which the Build Service-based 'KDE: KDE4:Factory:Desktop openSUSE_11.1' made more sense. Likewise, I added the OpenOffice.org and Firefox repos that had the latest versions of these apps; and of course, Packman, essentially for codecs.

After adding them, it's time to go back to the software overview section again. If you want to add your own custom RPM (or something that's third party), there's an option to upload that as well—refer to the 'Upload RPMs...' link in Figure 1.

The good thing about openSUSE is the availability of software patterns (similar to meta packages offered by other distros) that can help you install a set of programs together instead of individual apps one by one—thus making the job a tad easier. For example, to include most

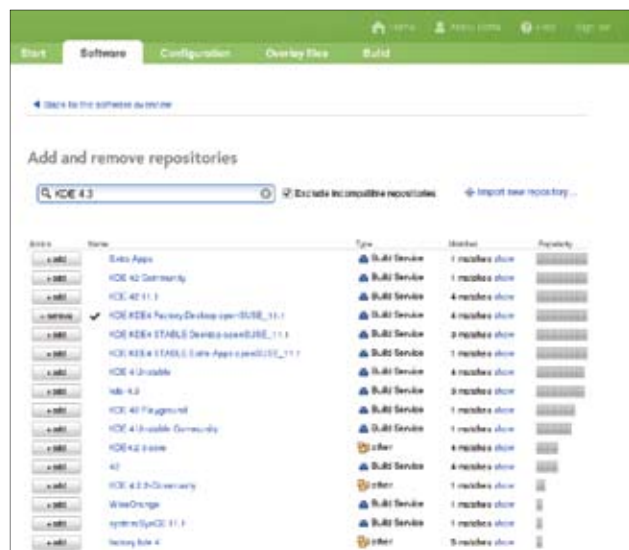


Figure 2: Search and add software repositories

of KDE 4.3 applications at one go, I first removed the *kde4* pattern that came from my base template, and added the one called 'KDE 4.3 Default Desktop'. Dependencies are all resolved and taken care of in the backend. When a package addition conflicts with another, it shows you the possible options to make amends.

The *Selected software* section is divided in two—patterns and packages. Note that the packages section will list all the packages that you've chosen (or automatically added when you included a pattern) but won't list the dependencies. This helps in keeping an eye on all the major packages that you select—and if required you can remove a package with a single click. The Web interface is based on the Ruby on Rails and Ajax technologies and that makes the service pretty interactive (a la Facebook, Gmail, etc).

Another of my requirements was to build an ISO image and keep the size within 700 MB so that I could pack it in a CD. Upon selecting a package, the left-hand pane of the Web page live-updates the download size, number of packages and patterns, besides there being an *Undo* button (or *Redo* when you delete a package) to keep things under check. It also shows you a warning if the size of the overall image becomes more than 700 MB. This is important because the 'Download size' info box on the left-hand pane is that of the individual packages that have been selected. A live ISO will have certain other things that take up more space. What I figured was that I needed to stop adding packages the moment the download size reached 630 MB, because the process of ISO creation adds another 70 odd MB to the overall size.

Once you're satisfied with the package selection, you can proceed for some finer-grained control over other settings, in order to personalise the appliance or distribution.

Get set

It's now time to access the *Configuration* section, which is further divided into seven sections.

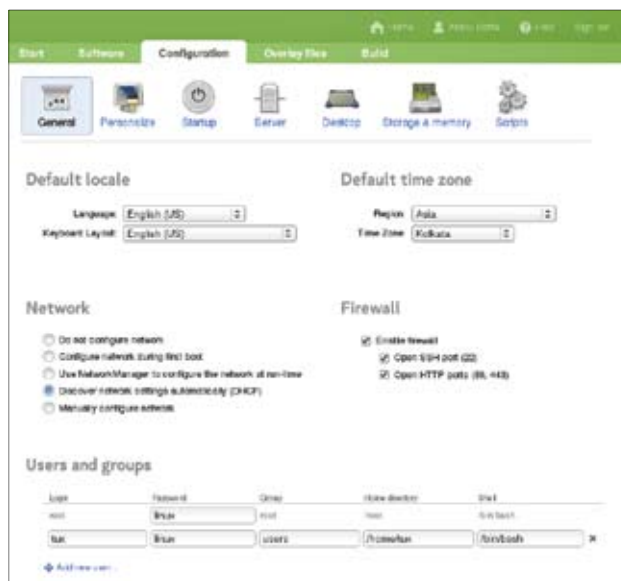


Figure 3: Configuring the general configuration options of our appliance

Under *General* (Figure 3), you'll need to set your language, keyboard layout, region and time zone choices. Next up, you can configure your network settings between an unconfigured network, configure on first boot, use NetworkManager, DHCP, and manual configuration. Leave the firewall off or on (with options to open default SSH and HTTP/HTTPS ports). And, finally, set the root password and add new users. By default, there's a user account called *tux* and its password (and that of the root's, too) is *linux*. You're free to change all these here.

In the *Personalise* section, you'll be able to set a logo and background for your appliance. There are a few images available by default with the option to add new logos and wallpapers. A nice touch by the Studio team is a live preview for Grub, the boot screen, the console and login screen option to give you an idea how they would look with your selection of graphics. I added a KDE logo to give it a more KDE feel and chose one of the wallpapers provided by SUSE.

Under the *Startup* tab you get to select the default run level (I wanted 5 for the graphical mode) and there's the option to add EULA (even multiple ones). The *Server* tab doesn't have much besides configuring the MySQL database, while the *Desktop* tab lets you set the user to auto-login and add start-up programs for users. The *Storage and Memory* section lets you set the RAM and disk size if you're creating a VMware or disk image. And, finally, the *Scripts* tab lets you add a custom script (if you wish to) at the end of the appliance build or when you finally boot it.

The next major section is called *Overlay files*. This one gives you a way to "...optionally add custom files to your appliance. Files added here will be copied into the appliance after packages are installed." You can upload files or set Studio to fetch



Figure 4: Building our appliance



Figure 5: Running a test drive by booting KDE 4.3 in the browser

it from a URL and then set the location to store the file, set the owner, group and permissions, etc. I didn't need this option, but it could come in real handy if you need it, for example, to upload a set of documents/media files, or even programs that the live user can access. There's another use that I'll come back to later.

Go!

Well, well, now that you're done with the configurations, it's time to hit the **Build** tab to run a build of the appliance/distro (see Figure 4). You'll require to select a build format between the live CD/DVD ISO, a disk image (for the hard drive and USB sticks), VMware image, and Xen image. We wanted a CD ISO, so if that's the case for you too, simply select that and hit **Build** and when done, download the ISO. By the way, the left-side pane is pretty smart. If you're creating an ISO image, it'll suggest you add the YaST Live Installer package as well, so that you can install the live CD to a hard disk if you wish.

But wait! The download takes too much time. Do we really need to wait that long just to check if everything is working as planned?

This is where the **Testdrive** feature (accessible when the build is ready) comes into play. You can live boot your appliance right inside your browser (see Figure 5) to test

Figure 6: After adding some modified files from *Testdrive* back to *Overlay files* section


whether everything is in place and working as planned. This also gives you another level of fine-grained customisation to add to your appliance using a feature called **Modified Files**. Note that this feature doesn't work when your build is an ISO, so choose any of the other build formats if you need to use this feature.

This tab lists whatever changes are done while you're running the image on your browser—every change is listed here since the **Testdrive** session began. So, if you want all the options and settings that you configured in your live session (for example, a set of active KDE widgets on the desktop) to stay in your final build, then go to this tab and select the changes that you need (while leaving out the rest that's not necessary) and add it to the appliance. Switch to the **Overlay Files** tab to see a new entry there (Figure 6). Read more about this feature at Nat Friedman's (SUSE Studio is his baby) blog: <http://nat.org/blog/2009/07/linux-in-the-browser>.

Kill the **Testdrive** session, rebuild the appliance and run **Testdrive** again to see if your changes appear fine, out-of-the-box. Once you're satisfied, go back to the build tab and this time build a final ISO, if you wanted that instead—like I did.

Now download it and test it on a few systems to see if everything works as planned. There is a possibility you might have missed out on something. Don't worry—go back to Studio, access your appliance and add, delete, build, **Testdrive** and download again—your appliance settings will be there till you delete them. Of course, you can create multiple appliances and keep all of them online for easy access and modifications at a later time.

Well, that's about it—to give you an idea of what's possible with SUSE Studio. Take it for a test run and check its potential.

It looks like building GNU/Linux distros and appliances of your own choice just got WAY too easy! Good job, Nat and the team at Novell! **END** 

By: Atanu Datta

He likes to head bang and play air guitar in his spare time. Oh, and he's also a part of the LFY Bureau.



The Perfect Desktop

A modern desktop with an office suite, multimedia codecs, wireless support and more. You can have all that in around 10-15 minutes.

*W*hat's stopping you from completely switching to a free software operating system? Multimedia codecs? Out-of-the-box support for the off-the-shelf hardware peripherals that you use? Games?

What else? (Send your 'what else' requirements to lfyedit@efyindia.com. We'd like to hear from you.)

Well, apart from support for a lot of games (the fault lies with the game developers anyway, that they don't have a GNU/Linux or BSD port of the game application), and certain niche/professional requirements, we believe a free software OS can provide you with everything—that too, in a user-friendly, fool-proof way.

Seriously, all of you (or your friends) who rave about the so-called out-of-box experience in Winduhs, try to install a Winduhs XP, Vista, or even the soon-to-be-released Winduhs 7, and look at what all you get.

Did I hear you say, "Only a handful of applications..."? And please tell me, honestly, if that Winduhs recognises all your hardware peripherals without you installing the individual drivers, one by one.

While you are reading this, I'd love it if you reset your PC/laptop BIOS to boot from a CD, and pop in this month's LFY CD into the drive.

Call it a coincidence, but the release of KDE 4.3 (a modern desktop environment) and the SUSE Studio service in August

enabled us to prepare for you what we believe has the potential to uproot that XP/Vista thing out of your system—only if you're willing to give it a try by learning and unlearning a few things. Give it a day (or at the most, a week), and if you don't shy away from asking questions, we believe you can make it.

Getting started: The first-timers' guide

Now that you've popped in the CD and reset your computer, you'll be presented with a menu along with some information. Just hit "Boot KDE 4.3 Live CD" to boot into a live environment—i.e., boot the OS.

For starters, a live CD runs directly off the CD by loading the programs into your RAM, without touching your hard disk in any way. The minimum requirements are 512 MB of RAM and a decent CPU that's not more than half a decade old. Newer and more powerful hardware would, of course, perform better.

You'll now see a bootsplash screen with a progress bar soon to be replaced by a licence agreement. Read it if you like—or just agree to it by hitting 'Yes' as most of us do anyway. Although, if you don't know the GPL, dedicating five minutes of your time to reading at least the first few paragraphs of the licence would make it more clear to you why free software is special.

After that, you'll go back to a black and white screen with some text scrolling. Don't panic—nothing is wrong here. You'll soon notice the graphical desktop start.

Yes, it's relatively empty—the only thing occupying the screen is a tiny panel at the bottom. Right click on the desktop and see what widgets you can run. I'd first recommend that you add the *Folder View* widget before anything else. This is like a container for your desktop folder, the place where you have all those icons, files and folders. The folder view comes in handy because even if you put your whole world on the desktop, everything is still accessible—nothing goes out of the screen because you get a scroll bar to help you navigate when this container fills up. Check out some other widgets if you like.

Hit the KDE menu (at the extreme left corner of the panel) to check all the applications we have for you. Here's a round up:

1. Office suite: OpenOffice.org 3.1
2. Browser: Mozilla Firefox 3.5.2 (besides Konqueror)
3. Audio player: Amarok 2.1.1
4. Video player: SMPlayer
5. Image editing: GIMP 2.6 and Krita 2.0
6. Digital camera tool: Digikam
7. Image viewer: Gwenview
8. Personal information manager: Kontact (mail, RSS reader, organiser, etc.)
9. Video editor: Kdenlive

One CD! It has it all (well, almost)

Well, let me share a little secret! Most of our regular readers would know that we generally bundle a KDE 4 Live CD (based on openSUSE) with the LFY issue that follows the specific KDE version release. The plan was the same this time.

The KDE Live CD has all sorts of pure KDE applications—but nothing else. For example, it doesn't have the OpenOffice.org (OOo) office suite nor Firefox. So, if you plan to install the CD, it means you'll have to download these two, for sure. Besides, most Linux distros don't bundle multimedia codecs out-of-the-box due to patent restrictions in some parts of the globe—just as with the KDE Live CD. Thankfully, in India, we don't have software patent laws! So, why should we be deprived, and unnecessarily spend on bandwidth to make the system usable?

I'll share another little secret: Pardus 2009 is perhaps one of the best drop-in replacements for your general desktop OS. It comes with the KDE 4.2.4, OOo 3.1 and Firefox 3.5.1, besides all sorts of multimedia codes. The only issue with it is that it was released before KDE 4.3 came out! So, we stole their concept, and aided by the release of SUSE Studio, we created our own Pardus replacement. Unfortunately, since none of us are developers here, our overall desktop integration is not as good as that of Pardus'.

10. Disc burning tool: K3b

Besides these, there's a chat client (Kopete), download manager (KGet), BitTorrent client (KTorrent), a PDF and document viewer (Okular), and various system and other miscellaneous tools.

There's support for compressing and uncompressing RAR, ZIP, LZMA, BZ2, GZ, etc, files aided by the Ark utility. Whether your music collection includes songs in MP3, ogg, wav and flac formats, Amarok should be able to manage all of them. SMPlayer will play most of the MPEG, DivX, XviD, MP4, flv, ogv, wmv, and other content you throw at it, apart from its ability to play encrypted DVD videos. Besides, you can play YouTube videos using Firefox.

I don't know what else you do, so please feel free to check the menu for applications. Here's a simple trick: hit Alt+F2 and in the pop-up box, fill in a term like, for example, 'audio' and see what application plays audio. Yes, you can launch that application straight from this tool.

To set your desktop properties, hit the KDE menu and launch *Configure Desktop*. You can finetune your desktop from there. Besides, the systems admin tool called YaST is where you can make system-level (as opposed to user-level) changes like configuring your network. The administrator (root) password is linux. So, if your system doesn't get its IP address

automatically from a DHCP router/modem (meaning that you need to fill in a static IP address for network access), or you need to enter your wireless access keys, you can do it here.

We have made sure that most of the Broadcom and Intel Wi-Fi cards should work out-of-the-box. In fact, if you're on a wireless network or need to frequently change your IP address, it's wiser to set up a tool called NetworkManager. Launch YaST, enter 'network' in the search bar and double click on the icon that reads *Network Settings*. You can 'Edit' your network settings in the *Overview* tab. To turn on the NetworkManager, switch to the Global Options tab and select the option that reads "User Controlled with NetworkManager". Then hit the *OK* button.

You should get a pop-up message near the system clock at the bottom right side of the screen that would report about a certain change. Locate an icon with an unplugged network wire in the system tray and configure your wireless/wired network there. The appearance of the icon should change to show a 'connecting', and then 'connected' status.



Tip: If by chance it fails after you're done with configuring the settings, log out (not restart) from the desktop and re-login as the live CD user 'tux' using the password linux, and try to connect to the network again. It should definitely work.

But this is only useful when the OS is on your hard drive and not on RAM, right? Which brings us to...

How about installing it on your hard drive?

Hit Alt+F2 and type 'live'. Yes, launch the *Live Installer* application to put the entire OS on your hard drive. There are two reasons for installing it:

1. Things will get a lot faster
2. Your changes will stay across reboots

Why? Because, right now everything is on the RAM—so it's heavily loaded—and the reason for everything seeming (sometimes, horribly) slow is lack of free memory. And, again, because everything is on the RAM, it all vanishes when you turn off your system.

The instructions on the installer are pretty simple and straightforward. The only place where you need some help is the partition screen on this installer wizard—i.e., only if you haven't installed a GNU/Linux system before. You should have at least 4 GB of free space on your hard drive where you can install the OS. Proceed only if you do. Else, free up some space using your preferred method.

On the partition screen, click on the button that says *Create Partition*; the next screen will show your hard disk, followed by a screen displaying your

partition table. Select the free partition and delete it! Now hit *Add* to create a new partition. Choose *Custom size* and enter 1GB there, before hitting *Next*. In the following screen, select the file system type as *Swap* from the drop down menu and click *Finish*. Again, hit *Add* to create another partition. Let it be the maximum size this time, and click *Next*. Let the file system type be ext3 now, and instead click on the option that reads "Mount partition". Enter / (i.e., slash from your keyboard) and hit *Finish*. That's it — you've successfully prepared your hard disk to install the OS.

This is, of course, not the best partition set-up. If you're interested to know more about partitions, go to <http://tldp.org/HOWTO/Partition/>

The rest of the installation wizard is simple. The time taken to copy the OS to your hard disk depends on your system speed. On a relatively new system with good hardware specs, it won't take more than 10 minutes. When it's all over, you should get a pop-up indicating the same and that you should reboot and then remove the CD from the CD tray.

Upon restarting your computer, the system will configure a few things and you'll land straight in the brand new desktop. There are a few more steps if you are one of those using the notorious Nvidia or ATI graphics chipsets. Follow the instructions at:

- <http://en.opensuse.org/NVIDIA>
- http://en.opensuse.org/Howto/ATI_Driver

...to set up your display drivers. Note that the version of openSUSE OS you're using is 11.1.



Tip: Your Firefox, GIMP image editor and OpenOffice.org suite will look ugly out-of-the-box (they are not native KDE applications). To make them look decent, launch the *Configure Desktop* utility, access the *Appearance* section and navigate to GTK Styles and Fonts. Set your style settings there—I prefer the *Clearlooks* style—and *Apply*. It'll prompt you to log out for the changes to take effect. Do that, and you're done.

While we're still here, I'd recommend you go to the *Fonts* settings section and change the fixed-width font from Monospace to Liberation Mono. This fixes the issue with the cursor position in Konsole.

So there you are; your free software operating system is up and running. Now, honestly tell me how long would it have taken you if you had wanted the same functionality on your freshly-installed Winduhs system. With that, I rest my case.

P.S.: You can read a review of KDE 4.3 on the following pages. 

By: Atanu Datta

He likes to head bang and play air guitar in his spare time. Oh, and he's also a part of the LFY Bureau.

KDE 4.3

It's the Desktop You Need, Almost



"KDE will very shortly become the desktop you need and not the desktop we think you need," was something stated by someone from the KDE team. Looking at 4.3, we think there's a fair chance that version 4.4 actually gets there.

What do you think of when you read something like this: "Fixed over 10,000 bugs and implemented almost 2,000 feature requests in the last 6 months; close to 63,000 changes were checked in by a little under 700 contributors..."? Now, that's what I call developers being really busy working on something solid.

Guess it's time to start our journey with KDE 4.3.

The first thing you notice in this version is the brand new theme called Air. This replaces the default and darker Oxygen that's been around since the inception of the KDE4 desktop. Air, by contrast, is a transparent white theme with nice stripes and spirals, which make it look quite appealing. Mind you, Oxygen is still available, besides various other themes that can be downloaded from KDE-Look.org by accessing the *New Theme* button from the *Desktop Settings* tool (you get it up by right clicking on an empty space on your desktop).

While we're on the subject of *Desktop Settings*—this tool has also received some attention. First, you have another type of activity setting—apart from *Folder View* (the traditional desktop with icons), and *Desktop* (the default desktop containment introduced in KDE 4.1)—dubbed '*Plain Desktop*'. It is basically a desktop without the toolbox you get on the top right corner to zoom-out of an activity and create new ones. Apparently, more activity types have been introduced, one of which can enable you to use your desktop screen as a whiteboard of sorts. However, I couldn't figure out a way to install these. [Let us know if you can.]

Another nice add-on under *Desktop Settings* is the additional wallpaper types—Virus, Mandelbrot, Weather and Globe, besides the slide show and patterns, which have been available for a while (Figure 1). Among these, I quite like the *Globe* add-on, which uses the *Marble* program to render



Figure 1: Wallpaper Types under Desktop Settings

an interactive glove as your desktop background.

Another pleasant surprise is that now you can have different activities for different workspaces. If you are the type who uses a lot of widgets, and finds it difficult to fit all of these into one screen, in earlier versions you probably needed to switch between activities to set some other widgets that wouldn't fit into your primary activity. However, let's face it, switching between activities is not the easiest of jobs. With version 4.3, you can configure each workspace as a separate activity and thus have different widgets, wallpapers, etc. on different workspaces (see Figure 3)—and thankfully, switching between workspaces is much easier.

However, the best feature among the 50-odd widgets that come with KDE 4.3, is the new enhancement added to the *Folder View* widget, which prompted me to switch to my home directory as my default location from the `~/Desktop` folder I have been using for a year. Now, hovering over a folder in *Folder View* lets me peek into the folder's contents, thanks to an overlay effect (Figure 2). With this sort of feature now available, you don't have to launch Dolphin to locate or access your files—just navigate the filesystem and access files straight from your desktop with a single (double) click.

To think about it, I must mention the inclusion of two widgets, the first of which is openDesktop. If you set the latitude and longitude on which you are located, this widget has the ability to show you other users



Figure 2: Peeking into the Folder View



Figure 3: Different activities for different workspaces

who are nearby. However, for this to work, you need an account with opendesktop.org; besides, there's not much you can do with the widget. It definitely shows some promise, and I won't be surprised if it becomes a major contestant soon among the various social networking tools.

The second widget seems handy for those who like to plan their days with different tasks, et al. It's called *Remember the Milk*, and its purpose is to keep you in sync with the RTM online service. Again, you need an online account to be able to use it. Besides these two, the Microblogging widget now has Identi.ca support, apart from Twitter.

The system notifications have

acquired some polish. Building on top of the notification system that was introduced in version 4.2, it now has a dedicated icon on the system tray. Running tasks now slide into this notification icon and are grouped together, not taking up too much of screen space. An animating icon signifies that jobs are still running. However, the 'More' and 'Less' buttons for notifications (for example, in the copy dialog), still need a lot of attention.

Coming to the various application enhancements, KRunner continues to improve. It now displays the search results as a list instead of as icons—the layout now gives it more room to display




Figure 4: New notification system



Figure 5: The new KRunner layout

the application descriptions below their names. The option to run the apps as a different user is also back, which was available in the KDE3 series, but went missing in KDE4. Also, the addition of the help button (the ? icon) gives the user an idea of the different syntaxes that we can use. Earlier, it more or less depended on guesswork.

The Dolphin file manager now shows a slideshow of the images that are inside when you hover over a folder when using the file 'Preview' mode. A feature to play (preview) media files in the 'Information' pane has also been added.

The optional tree-view is back in *System Settings*, reminiscent of KControl from KDE3, for those who need it. A major bug fix, I'd say, is the ability to configure items that require root-level changes (like, for example, KDM) from within *System Settings*. Any option that required root privileges used to be greyed out earlier—now it prompts for a root password instead. Also, PolicyKit has been integrated into the desktop and a module is made available here. With this you can fine-tune certain privileges. Like, for example, Dolphin usually lists all available partitions in the side pane, whether an entry has been set in the */etc/fstab* file or not. However, in previous KDE versions,

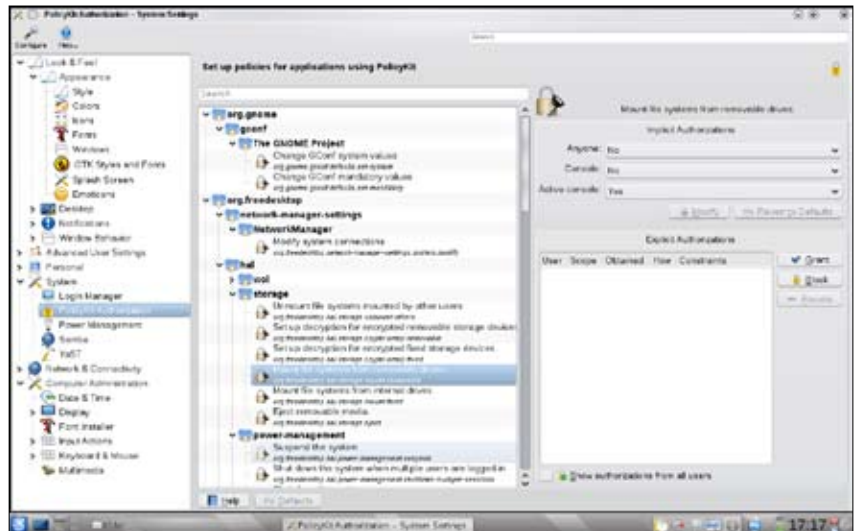


Figure 6: System Settings using the Classic Tree View layout (see left pane); fine tuning PolicyKit settings (right)

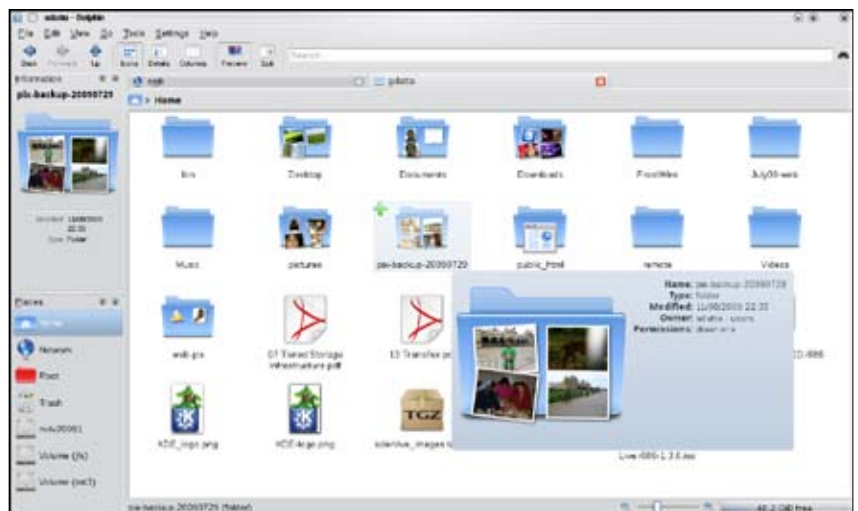



Figure 7: Image preview in Dolphin file manager

trying to access such a partition would result in a *Not permitted* error message. Now if you set the settings correctly in the PolicyKit module, you can give a normal user the privilege to mount these partitions.

All these are small add-ons, which makes 4.3 more of an incremental update rather than a major overhaul like its predecessors. There are, of course, numerous other enhancements; I've only listed the ones that drew my attention.

The most impressive factor is the overall stability of the applications and Plasma, and thereby the stability of the desktop as a whole. However good KDE4 has been till date, it's always been a tad unstable

compared to its predecessor or alternative desktop environments. Having used 4.3 since the RC1 release, I've hardly faced crashes. And when it does crash, the crash report tool has a button to restart whatever has crashed—provided it's a KDE app.


The news is that work had already started on version 4.4 much before 4.3 was released; so expect many new additions in the version that comes out next January. **END** 

By: Atanu Datta

He likes to head bang and play air guitar in his spare time. Oh, and he's also a part of the LFY Bureau.

Who Says Linux Is Not (in) Rocket Science?

People have always looked at the sky and dreamed of flying to the stars. However, throughout history, it was impossible to create a reliable engine that could overcome the Earth's gravity and pull a payload into orbit. Luckily, the scientific achievements in the 20th century—in chemistry, in physics and in the theory of interplanetary flights—have allowed mankind to build rocket engines capable of not only carrying a satellite, but also humans.



Perhaps we all remember the time when the first artificial satellite, Sputnik, was launched into a low altitude elliptical orbit around Earth. That was October 4, 1957. In general, you can look at this date as the beginning of an era of human space flights. Sputnik [1] was a very simple device by today's standards. It could only transmit radio signals («beep-beep-beep») at frequencies of 20 and 40 KHz, besides being only 58 centimetres in diameter and weighing around 800 kilograms—a very tiny unit compared to its carrier rocket, R7, which had a height of about 30 metres and weighed 280 tons.

The history of rocket development has roots in ancient times, when gunpowder was invented and when the first fireworks appeared in China. We can't say for sure whether it was clear back then how to create a missile and why, if at all, it ought to be created. However, this firework became the basic element of a cannon—in other words, modern artillery. Cannons used a controlled explosion of gunpowder—the energy that was released pushed a missile towards a certain direction. Afterwards, a whole new science in artillery came up, called ballistics—another fundamental element that helped launch rockets.

Sergei Korolev [2] and Werner von Braun [3, 4] were born in the beginning of the 20th century, and both went on to become the fathers of space flight. von Braun developed the legendary V-2 rocket [5], which was really a predecessor to the Russian R7 (a rocket that pulled the very first Sputnik into the sky and set it into an orbit around Earth) and was responsible for the American Saturn V (another great rocket that helped humans reach the moon in 1969).

I must say that rocket design and development is a very long and expensive process. For example, the first real implementation of rockets started in the early 1930s in Germany. In 1934, the first V-2 (A1) rocket prototype was successfully launched. It was then named the Aggregat (thus, A) rocket. From 1941 till 1945, there were more than 200 test runs of the V-2.

In order to build the successor of V-2, the R7, the Russians created more than a dozen research institutes, with more than hundreds of scientists working there.

But all this was half a century ago. What are we up to now? Can we dream about an amateur rocket that is several times smaller than the V-2, R7 or Saturn V, and can shoot up at least 10 kilometres? Yes, we can! And this is a reality—an amateur rocket with Linux as a core system!

A rocket as a student graduate project!

In 1997, two undergraduate electrical engineering (EE) students from Portland State University, Oregon (USA)—Andrew Greenberg and Brian O’Neel—had a very simple plan for their graduate project — they wanted to create one that was: a) challenging; b) fun; and c) potentially hazardous for them, possibly even fatal! Andrew recalls, “Plugging components into breadboards in student labs didn’t fit the bill. We realised that what we really wanted to do was to put something into orbit ourselves.” And that’s what they decided to do: build a rocket that



Figure 1: Werner von Braun near Saturn V

would be capable of inserting a 1-kilogram pico-satellite into orbit.

With this dream, Greenberg and O’Neel established the Portland State Aerospace Society (PSAS). During initial research, they found out that while large amateur rockets existed, they were all ballistic rockets—which meant, they shot up, coasted for a while, and then fell down. Andrew recalls, “No one had yet built an amateur rocket that was capable of following an orbital insertion trajectory. So that’s the road we started down: building a medium-sized rocket that used cheap, off-the-shelf parts with open source software to create an orbital-capable launch vehicle.”

The very first prototype LV0 (Launch Vehicle) was really just a “how does all this rocket stuff work anyhow” kind of a rocket. “We purchased the airframe from a hobby store, and shoved our own electronics system in it: a small microcontroller with an accelerometer and pressure sensor. Glenn LeBrasseur, another EE student back then, was an amateur radio enthusiast. So we threw in some radios broadcasting data and video. Shockingly, it mostly worked.”

LV1 was what you’d call a second-system syndrome—it was quickly built, not very carefully thought out, and was overly complicated. However, as Greenberg shares, “It



Figure 2: The Russian R7 launch rocket vehicle for Sputnik

was a good experience because it was our first real avionics system: we had an 8-bit microcontroller that processed a six degree of freedom inertial measurement unit (three accelerometers and three gyroscopes), a GPS receiver, pressure sensors, and could fire the pyrotechnic charge to release its parachutes. Again, shockingly, it worked fairly well.”

However, after the LV1 launch, it became clear to them that the 8-bit microcontroller was way too underpowered: “Writing in assembly language was ridiculous, and primary batteries that lasted 20 minutes was a bad idea. Besides, there were a lot of issues with LV1.”

With LV2 rockets, they did the right thing—they wrote requirements, did research on systems architecture and design, and got other engineering students from the computer science and mechanical engineering faculties involved. Then those at PSAS thought of a very smart idea: what if they relied on open source in the form of GNU/Linux and open source tools to solve some of the requirements... would that help them build their rocket? With this came a server infrastructure running wikis, mailing lists, repositories for code and CAD, besides a legal framework, like the GPL.




Figure 3: All modern rockets share the same concept as V-2, but instead of a warhead, the amateur rocket carries a useful payload

“And so LV2 ended up with a 32-bit processor running Debian Linux 2.4. Our launch tower became another Linux box, and we had a field server routing packets between the rocket, launch control and the launch tower... it was all fantastically fun. And for the first time, we really had enough horsepower in LV2 to consider starting active guidance,” reveals Greenberg.

Open source components

Currently, to build and launch the rocket requires the combined efforts of several teams with their respective specialisation in airframes, avionics, communications, ground control, the payload, propulsion, software and uncertainty. As you can see, this is a mixture of software and hardware teams—and they all play important roles.

The airframe team is responsible for designing the LV rockets—how to construct an aluminium module, where to place motor casing, nozzles and coupler bulkhead, as well as the aeroshell construction and parachute recovery module, which hides inside a rocket and is used on its return to the ground.

The avionics team produces all the software [6] and deals with hardware modules (like the flight computer, discussed below). This team is responsible for reliable software interconnections during flight and after-flight analysis (via the syslog Linux-subsystem). For example, the team is deeply connected with the ground team in terms of radio-exchange and communication, because all radio links from a rocket to the ground (amateur TV, Wi-Fi connections) are managed by the

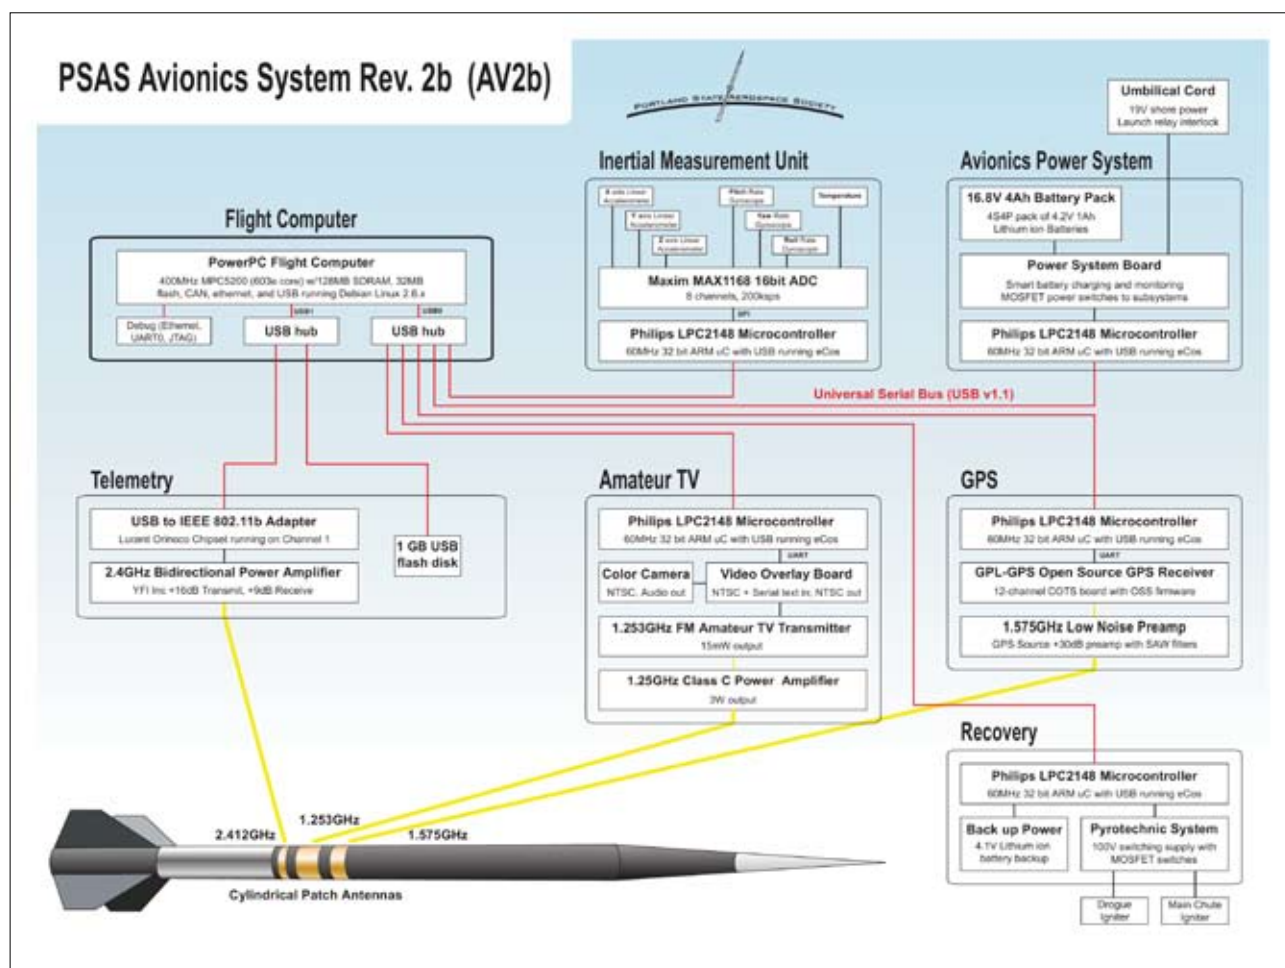


Figure 4: The avionics and flight control system is based on Linux

ground team. Pre-flight and flight videos [7-10] are recorded by software made by these teams.

Another amazing group is the propulsion team that conducts research on different chemicals like what propellant ought to be used for rocket engines—LOX/Paraffin (liquid oxygen, plus a kind of wax), GOX/Paraffin (gaseous oxygen), or a very standard fuel like ammonium perchlorate [11] oxidiser with hydroxyl-terminated polybutadiene (HTPB) fuel.

Those from different teams very often form a new team for closer cooperation, as is the case with the payload team. Its members are from the avionics and airframe teams.

However, the most interesting perhaps is the Uncertainty Team! Its task is to achieve things that are unpredictable and very challenging. This team does research and implementation of probabilistic algorithms, and in particular, all sorts of Kalman filtering (a mathematical technique widely used in control systems and avionics to extract a signal from a series of incomplete and noisy measurements): Kalman unscented, Kalman extended, sigma point Kalman, particle filters and others. These are useful in estimating the position and orientation of a rocket given a diverse array of sensor inputs, all of which are incorrect in their own way. Without such estimation and the eventual feed back of this information to the whole system, controlling a rocket's trajectory can be extremely difficult.

Discussing all these details is out of the scope of this article; however, we're certainly concerned about the open source software parts. So let's take a closer look at the operating system used by the avionics team.

The current flight computer for the team's rocket is a PowerPC-based single-board computer that was specifically chosen because of a "Linux on Power" grant the students managed to win from IBM. It's a 400 MHz Freescale MPC5200 PowerPC

single-board computer (SBC) from TQ Components. The following are its specs:

- CPU: 400 MHz Freescale with a 603e core
- RAM: 64 MB SDRAM/32 MB Flash
- Peripherals: USB 1.1, CAN, UARTS, PCI, ATA/IDE, SPI, etc
- Size: 80 x 60 mm

This flight computer uses U-Boot as the kernel boot-loader, a kernel image itself, then a device tree blob with an initramfs disk image that is stored in Flash. When the flight computer boots, the bootloader loads everything into the RAM and then boots from it. This means that all changes to the filesystem will be wiped out on the next power cycle. That's why the system design has the following features: a) all software you ever want to use during a flight is installed in Flash; and b) rocket logs are stored on a compact Flash card to avoid being lost by an in-flight power cycle.

The kernel used for this PowerPC board, is stored at Denx's site (developer and maker of TQM5200 board) and can be easily grabbed with git:

`$ git clone git://git.denx.de/linux-2.6-denx.git`

All other operations required for the flight computer are: Linux kernel compilation and U-Boot loader with a cross-compilation suite found in the Debian package repository, building user-space programs and drivers (including CAN and USB), making the initramfs image, and combining it all together to fit on a 32-MB Flash card.

Other major components used during the launch and flight are the telemetry unit, the GPS and recovery unit, as well as the inertial measurement unit, the power system unit and the amateur TV unit. All of these are based on the LPC2148 board with an ARM processor:

- CPU: NXP's ARM7TDMI LPC2148 with integrated USB 2.0 device
- RAM: 512 KB Flash / 32+8 KB



Figure 5: The small, yet very powerful PowerPC flight computer



Figure 6: RocketView is the main application that doesn't leave the ground

SDRAM

- Clock frequency: 12.0000 MHz (equals 60 MHz CPU clock)
- Size: 59 x 39.5 mm

The operating system used in this tiny board is eCos [12]—a real-time operating system intended for embedded applications. Rocket developers have a dedicated git-tree, because of a lpc2148-patch:

`$ git-clone http://psas.pdx.edu/~sarah/ecos.git/`

Cross-compilation and building for the LPC2148 board is similar to the process used for the flight computer.

Once all the software components are built and written to the Flash card, it is time for the hardware teams to put it all together.

But for now, let's look at how to create a PCB. There has been some debate on what CAD software ought to be used: KiCAD, gEDA or Eagle. "We look forward to using open source CAD tools for electrical and mechanical design, but we don't think any of them are quite there yet. Some are very, very close, and we hope to help them out in the near future, but for now we're using closed source electronic and mechanical CAD tools," says Greenberg. So, it

was decided that Eagle is much better for the design. Though it isn't an open source tool, it is free for non-commercial/academic applications and pretty much multi-platform.

Next step? A control software that is left on the ground—RocketView runs on a laptop, and is operated by the launching staff. This telemetry display software was originally written with GTK+ for visualising those events that come from the rocket via the radio link. Later, it was rewritten as a Java application with the same functionality, plus launch control functions. You can now safely push a button and a rocket will fly into the sky.

Fuel and further development

So what is the fuel that boosts the rocket engine? According to Sarah Sharp: "It's ammonium perchlorate." Greenberg elaborates, "We did indeed have a GOX/Paraffin engine project, but it's been on hold for several years now. We hope to restart that project sometime in the future. For now, we're sticking with very standard solid amateur rocketry motors: an ammonium perchlorate oxidizer with HTPB fuel. We launched on a small "N" motor this last May, and we're aiming for a larger "N" this October at the Black Rock Desert




Figure 7: The LV2 rocket on the tower, moments before launch



Figure 8: The PSAS community with its LV2 rocket

amateur rocketry launch. [13]"

However, there is great potential in GOX/Paraffin and LOX/Paraffin technology. While it is absolutely safe to be used in rocketry, it is also much cheaper compared to any alternative. Besides, there is big research work on, which will allow hybrid motors with Active Fin Control (AFC), Thrust Vector Control (TVC), as well as a reaction control system (RCS) to use it. It could enable achieving a new flight record that's above the already-achieved altitude of 10 kilometres.

Greenberg reveals, "PSAS is going to reach higher altitudes sometime later next year. We're currently busy with rebuilding the avionics system, so you can subscribe to a mailing-list and help us reach the sky." 

Acknowledgement

The author would like to thank the following PSAS members for their help during this survey: Andrew Greenberg, Sarah Sharp, Jamey Sharp and Ian Osgood. Pictures are copyrights of Sarah Sharp, Wikipedia and its members.

Resources

- [1] http://en.wikipedia.org/wiki/Sputnik_1
- [2] http://en.wikipedia.org/wiki/Sergey_Korolyov
- [3] http://en.wikipedia.org/wiki/Wernher_von_Braun
- [4] <http://history.msfc.nasa.gov/vonbraun/bio.html>
- [5] <http://en.wikipedia.org/wiki/V-2>
- [6] <http://psas.pdx.edu/CapstoneLV2bProjectReport/>
- [7] http://psas.pdx.edu/news/2009-05-31/2009-05-31_helicopter_overflight.mov
- [8] http://psas.pdx.edu/news/2009-05-31/launch_dk.avi
- [9] http://psas.pdx.edu/news/2009-05-31/launch_600fps.mov
- [10] http://psas.pdx.edu/news/2009-05-31/launch_onboard.avi
- [11] http://en.wikipedia.org/wiki/Composite_propellant
- [12] <http://ecos.sourceforge.org/>
- [13] <http://www.balls18.com/>
- [14] http://commons.wikimedia.org/wiki/File:Semyorka_Rocket_R7_by_Sergei_Korolyov_in_VDNH_Ostankino_RAF0540.jpg
- [15] http://commons.wikimedia.org/wiki/File:S-IC_engines_and_Von_Braun.jpg
- [16] http://psas.pdx.edu/news/2009-05-31/2009-05-31_helicopter_overflight.mov
- [17] http://psas.pdx.edu/news/2009-05-31/launch_dk.avi
- [18] http://psas.pdx.edu/news/2009-05-31/launch_600fps.mov
- [19] http://psas.pdx.edu/news/2009-05-31/launch_onboard.avi

By: Anton Borisov

The author has specialised in Linux and FOSS technologies for more than a decade. His professional spheres of interest include, but are not limited to, robotics, embedded systems, statistics and algorithmic methods.

The Gol is Just Not That Into



The ministries and babus of the Indian government departments are all heading towards the road to slavery.

Different agencies of the Indian government seem to be making all the wrong moves lately. Two issues have come to our notice and both are quite disturbing. These issues not only raise questions over how such national bodies discriminate against free software alternatives, but also reveal how they are forcing people to use non-free or proprietary software sold by big corporations.

The irony is, while one government organisation endorses proprietary software, the other releases software that is even more restrictive than those sold by their commercial counterparts. The two organisations in question are Technology Development for Indian Languages (TDIL), established by the Department of Information Technology (DIT), and the Indian Space Research Organisation (ISRO).

Fonts: Can't read the free word

The apex body of the DIT recently finalised the fonts to be used for e-governance related work. While the apex body approved Unicode 5.1.0 as a standard for e-governance applications for all 22 Indian languages, except for Kashmiri, it chose a font that is not free software compliant.

The font approved is Sakal Bharti, which has been jointly developed by TDIL and CDAC. But the font itself is not 'free' in nature. If you go to the TDIL site to download the font approved by the apex body, you will be greeted by a disclaimer that forces you to agree to certain terms if you want to download it. I, as an individual or a company, ought to have the full freedom to download and use anything that has been developed by using taxpayers' money. The owners of such technologies should be the citizens

who pay taxes that aid in the development of such fonts, and not the body that developed it, using taxpayers' money.

So, if I go by their licence agreement when I want to download a font, I will have to agree to: "1. These Products are for Academic/R&D/Personnel use only." Then another restrictive clause: "3. Copies of the Software and the manual can be made only for back-up purposes and non-commercial personal products/utilities/fonts usage. Any copy made must include all copyright and proprietary information notices appearing on the copy provided herein. Such a copy, if later on used for commercial purposes, will be treated as illegal."

What if I want to use these fonts for my commercial website? Not being permitted to do so, according to me, is against the principles of the free software philosophy.

According to the FSF, "Free software does not mean non-commercial. A free program must be available for commercial use, commercial development, and commercial distribution. Commercial development of free software is no longer unusual; such free commercial software is very important. You may have paid money to get copies of free software, or you may have obtained copies at no charge. But regardless of how you got your copies, you always have the freedom to copy and change the software, even to sell copies." [www.gnu.org/philosophy/selling.html]

But all the conditions for the 'government approved' font are quite contradictory to the FSF's philosophy.

(Editorial team's note: Even if you want to disregard the principles of free software and what the FSF preaches, I guess we agree that if it's developed with the aid of funds from the tax payers' money, it should be in the public domain. Period!)

This may lead to many more complex situations as this font is not at all compatible with the GPL, or any of the other free software, for that matter. Thus, even the government bodies won't be able to use these fonts with a lot of free and open source software (FOSS) they apparently deploy and implement.

When asked how appropriate the move was and what would be the ideal situation for choosing fonts for official work, Richard Stallman said, "It is flagrant injustice for the government to distribute useful information (such as software or fonts) to the public with these restrictions."

He added, "We tried for years to get the CDAC fonts to be freed, but I think they never did so." In the end, he warns that, "Fonts are works of practical functional use, so their users deserve the four freedoms. Furthermore, fonts today normally take the form of software, so a non-free font is non-free software."

I am curious; when the font does not comply with

the free software philosophy and slaps restrictions on users, then: (1) How can those officials sitting in the apex bodies choose it, knowing it is not truly 'free' in nature; and (2) How can an organisation like TDIL (which is supposed to serve the public interest) release software under such restrictive terms and conditions — which they call a licence.

The reason I've honed in on TDIL is that it is a programme started by the Department of Information Technology (DIT), with an aim to develop information processing tools to facilitate human machine interaction in Indian languages, and to develop technologies to access multi-lingual knowledge resources.

It is also debatable if the Union government is even concerned about free software at all. It was the opposition parties who argued against the software patent deal—and the CPI(M) took the lead in that. Also, the BJP, though infamous for its religious dogma, supports free and open source in a big way by going as far as building their whole party IT infrastructure on FOSS. It's a shame that the recent activities by government bodies shun the opportunities of freedom—or do they say no to FOSS only because the opposition had said "Yes" to it?

Bhuvan: Continue to pay *lagaan* to Microsoft

ISRO recently released Bhuvan and the entire nation has gone gaga over it. However, ISRO seems to believe that there are no other operating systems besides Windows—or do they think that this is all Indians are capable of using? This is ironic, considering that the government-funded CDAC has its own GNU/Linux operating system called BOSS. So, it appears that CDAC is being BOSSed around by proprietary and non-free counterparts.

According to an ISRO press release, "The Bhuvan geoportal enables users to access information on basic natural resources in the geospatial domain, particularly the Indian images and thematic information in multiple spatial resolutions. This would provide a sharper picture of Indian terrain barring sensitive locations such as military and nuclear installations."

It also says that there would be many more value-added functions and facilities that will be added into the package from time to time. "A particular interest of ISRO/DOS would be to provide such functionalities to the common man so that he/she adopts a participatory approach with scientists to solve simple problems easily and interactively."

But, how did ISRO figure out that the common man of India uses only Microsoft Windows, because if you go to the Bhuvan site, you will be welcomed with a garland that says, "View in IE 6.0 or above only with 1280 x 1024 resolution."

And here're the system requirements to run Bhuvan:

- Operating system: Windows XP/Vista

- RAM: 512MB
- Hard disk: 2GB free space
- Network speed: 256 Kbits/sec
- Graphics card: 3D-capable with 32 MB of VRAM
- Screen: 1280 x 1024
- 32-bit True Colour


But wait... there's more: "To browse Bhuvan, you require the Bhuvan plug-in, which can be downloaded from this website, after registration and you will also need DirectX 9.0 or a higher version (www.microsoft.com/windows/directx/) and the MS .NET framework 2.0 or above for installing the plug-in. Please note that the Bhuvan plug-in can be installed with administrative privileges only."

Now, while they call this a 'Google Earth killer', most people can't even see if it works or not, because many don't use Windows—at all. And why should they waste money on buying a licence for an operating system just so they can run Bhuvan, when they are perfectly happy with the current OS they are using? Well, thank you for your offer GoI and ISRO, but we're perfectly okay with Google Earth—although non-free software, it still runs on GNU/Linux. And hey, the free software alternative called Marble [<http://edu.kde.org/marble>] is also out there. Why couldn't ISRO work with them, to collaborate and share, all for the greater good that benefits all?

And what's with the only-runs-on-Internet Explorer

factor? Yes, all those who're still stuck with Windows, yet have switched to Firefox or something else for their browsing needs are now required to switch back to IE if they want to use it. Seems like ISRO is doing free-of-cost PR and sales work for Microsoft (well, IE's browser market share is at an all-time low, and seems to be dipping every day) with the Indian taxpayer's money.

These Indian government bodies seem to be going where water doesn't flow freely. This is discouraging for our economy, that too for a sovereign country like India. Governments don't wake up till it's election time again. Seems like our only hope remains with the opposition parties (besides the free software community), to not let the government sell us out to proprietary corporations that are not even based on Indian soil.

A copy of the minutes of the apex body meeting is available at: http://fosscomm.in/OpenStandards/Apex_Body_Meeting_Minutes **END** 

By Swapnil Bhartiya

A Free Software fund-a-mental-ist and Charles Bukowski fan, Swapnil also writes fiction and tries to find cracks in a proprietary company's 'paper armours'. He is a big movie buff, and prefers listening to music at such loud volumes that he's gone partially deaf when it comes to identifying anything positive about proprietary companies. Oh, and he is also the assistant editor of *EFYTimes.com*.

LINUX

ForYou

Calling...

Share your knowledge

Inviting FOSS experts to write articles on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of www.openITis.com (aka linuxforu.com) we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at lfyedit@efyindia.com

Few topics that top our list:

- ◆ Tips 'n' Tricks for software developers or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools





Play With Your Partitions

This article explains some free and open source tools like GParted, g4u and partimage, which are used for partitioning, disk cloning, partition backup/restore, etc.

Let's face it—hard disk-related activities can never be simple, especially when that disk has data on it. In this article, we'll take a look at some of the tools that can make these activities easier.

A partitioning tool

There have been many operating systems and live CDs that provide partitioning utilities. Often, they turn out to be either proprietary or not able to support all filesystems used for Linux, Solaris, etc. But GParted is one such important partitioning utility that supports many filesystems and comes with a friendly user interface.

GParted (Figure 1) stands for 'GNOME partition editor'. With it, you can create, modify and delete partitions. Although you can definitely install it on any distribution and run it, I wouldn't recommend running it on a booted OS as problems might occur with certain mounted partitions. A better solution is to run GParted from a live CD. GParted itself comes on a separate live CD; boot it and double click on the GParted icon from the UI presented. It is also a part of a live CD called System Rescue CD.

To run GParted from here, first start X by issuing the *startx* command, and then open the GParted application from the X terminal. It can also be run from a bootable USB drive or via a PXE boot.

Once GParted is run, you can use it for the following operations:

- To create a new partition table (by erasing the old one, if one exists). Use this option very carefully. Unlike other operations, *Undo* is not supported for this.
- A new partition can be created from unallocated/free space. Decide on new partitions as primary, extended or logical, as required.
- To delete an existing partition.
- To resize/move existing partitions (Figure 2). Sizes can be increased only when free space is available on any adjacent side.
- To change flags/labels of a partition.
- To format a partition to any supported filesystem—ext3, reiserfs, etc.
- To check filesystems for errors.
- To view partition information.

All requested operations are added to the pending queue, which is performed in sequence upon clicking the *Apply* option.

The following filesystems are supported

by GParted presently: btrfs, crypt-luks, ext2, ext3, ext4, fat16, fat32, hfs, hfs+, jfs, lvm2 pv, ntfs, reiser4, reiserfs, swap, ufs and xfs. Check GParted.sourceforge.net/features.php for individual support of each filesystem and for updates.

Disk cloning

g4u is a disk cloning software. It is helpful in larger organisations like educational institutes where there are several machines of identical configuration. Every system may need multiple operating systems along with some updates and customisation. Repeating these for every system is a tedious task. So cloning is a good technique to avoid such problems by copying many disks from one fully prepared one. Also, when one or more operating systems crash on one system, it can be recovered quickly by copying the OS from an identical system.

This software has a few advantages over other cloning utilities. One is that it's free. Second, it is independent of partitions, filesystems and operating systems that exist on a disk, as cloning is performed block by block. The only requirement is that the target disk should be larger than the original disk.

g4u comes as live media, in the form of two bootable floppies or a bootable CD. The images for floppies/CD can be downloaded from www.feyrer.de/g4u/#reqs. As of today, the latest stable version of g4u is 2.3.

Disk cloning can be done in two ways — locally, by connecting two or more disks to a system at a time, and remotely, where the image of the original disk is stored on an FTP server, which is used for other disks to be copied using FTP.

So, to perform disk cloning, connect two or more disks for the local copy, or one disk and the network set up if you plan to use FTP, and run g4u from a live CD or bootable floppies. A terminal will be presented to use the following commands:

1. **disks** lists all connected disks, and specifies identifiers for each. IDE hard disks are identified by wd0, wd1, etc, whereas SATA disks are identified by sd0, sd1, etc. Here, wd0/sd0 represents the first/primary hard disk.
2. **parts <disk-name>** lists the available partitions in a specified disk. For example, **parts wd0** lists partitions of disks identified by wd0. See Figure 3.
3. **copydisk <disk1> <disk2>** copies all data from disk1 to disk2, block wise (Figure 4). For example, **copydisk wd0 wd1**.
4. **copypart <d1p1> <d2p2>** copies a partition from one disk to a partition on another disk. For example, **copypart wd0e wd1f**, copies the second partition of disk1 to the third partition of disk2. (Here the letters d, e, f, etc, in no way concern the drive letters of any OS. g4u assigns a character for each partition, starting from d for the first partition.)
5. **uploaddisk** stores the image of the whole disk on an FTP server:




Figure 1: A view of GParted



Figure 2: Resize and move a partition

```
g4u> parts wd0
Found Linux Ext2 partition: size 41945652 (20481 MB), offset 63
skipping existing Linux Ext2 partition at slot e.
Found Linux Ext2 partition: size 62910477 (30718 MB), offset 41945778
skipping existing Linux Ext2 partition at slot i.
Found Linux Ext2 partition: size 57335922 (27996 MB), offset 104856318
skipping existing Linux Ext2 partition at slot j.
Found swap partition: size 5574492 (2721 MB), offset 162192303
skipping existing swap partition at slot k.

11 partitions:
#    size    offset    fstype [fsize bsize cpg/sgs]
d: 167772160 0        unused 0 0          # (Cyl. 0 - 166440*)
e: 41945652 63      Linux Ext2 0 0          # (Cyl. 0 - 41612*)
i: 62910477 41945778 Linux Ext2 0 0          # (Cyl. 41612 - 104024*)
j: 57335922 104856318 Linux Ext2 0 0          # (Cyl. 104024 - 166904*)
k: 5574492 162192303 swap      0 0          # (Cyl. 160905 - 166435*)

Not updating disk label.
g4u>
```

Figure 3: Output of the **parts** command

uploaddisk ftp-address imagename disk-id

...where **ftp-address** is the IP address of the FTP server. By default, the FTP account named 'install' is used. Use **account@ftp-address** for any other account. **disk-id** is the identifier of the hard disk (for example, wd0, wd1, sd0, sd1, etc) whose image needs to be uploaded. It is optional. The absence of **disk-id** implies taking the image of the first hard disk, which is identified by **rwd0d**. Also, by default, gzip-9 compression is used while saving images with a typical extension like .gz. To use lower levels of gzip compressions (higher levels of compression take a longer time to back up/restore) use the following:

GZIP=1 uploaddisk ftp-address imagename disk-id

6. **uploadpart** stores the image of a particular partition on the FTP server:

uploadpart ftp-address imagename disk+part

```
* This screen: help

[disk] defaults to wd0 for first IDE disk, [disk+part] defaults to wd0d for
the whole first IDE disk. Use wd1 for second IDE disk, sd0 for first SCSI
disk, etc. Default image for slurpdisk is 'rud0d.gz'.

Enjoy!                               Send comments to hubert@feyrer.de
                                     Donate at paypal@feyrer.de!
                                     http://www.feyrer.de/g4u/

g4u> disks
wd0 at atabus0 drive 0: <UBOX HARDDISK>
wd0: drive supports 128-sector PIO transfers, LBA addressing
wd0: 81920 MB, 166440 cyl, 16 head, 63 sec, 512 bytes/sect x 167772160 sectors
wd0: 32-bit data port
wd0: drive supports PIO mode 4, DMA mode 2, Ultra-DMA mode 6 (Ultra/133)
wd1 at atabus0 drive 1: <UBOX HARDDISK>
wd1: drive supports 128-sector PIO transfers, LBA addressing
wd1: 81920 MB, 166440 cyl, 16 head, 63 sec, 512 bytes/sect x 167772160 sectors
wd1: 32-bit data port
wd1: drive supports PIO mode 4, DMA mode 2, Ultra-DMA mode 6 (Ultra/133)
wd0(piixide0:0:0): using PIO mode 4, Ultra-DMA mode 2 (Ultra/33) (using DMA)
wd1(piixide0:0:1): using PIO mode 4, Ultra-DMA mode 2 (Ultra/33) (using DMA)
g4u> copydisk wd0 wd1
3441 KB 264.42 KB/s
```

Figure 4: Disk cloning in progress

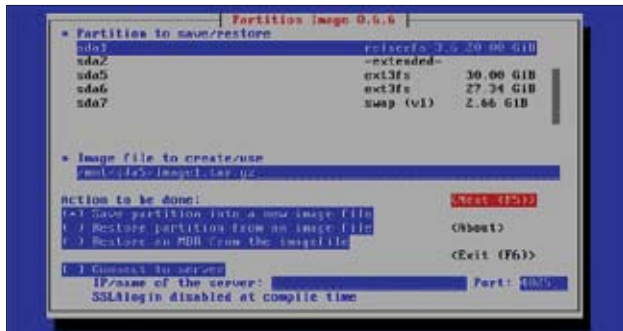


Figure 5: Initial screen of partimage

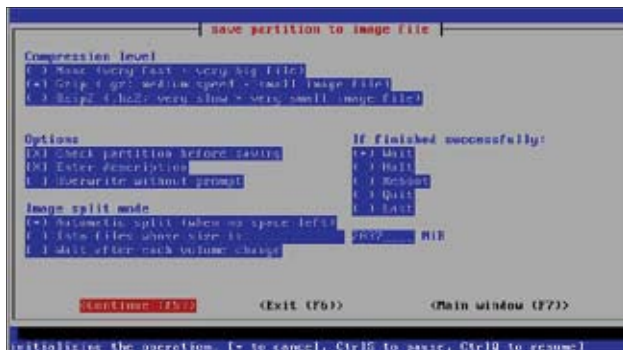


Figure 6: Options to save the partition

7. **slurpdisk** restores a disk from the FTP image:
`slurpdisk ftp-address imagename`
 8. **slurppart** restores a partition from the FTP image:
`slurppart ftp-address imagename disk+part`
- It is preferable to run g4u in a DHCP-enabled network if FTP transactions are planned. If DHCP is not available, the network needs to be set up manually.

Partimage for back up/restore

Consider the following scenario. You have installed a Linux distribution that comes with minimal packages. Additional packages have been installed from the Internet using a package manager. So, patiently, you've updated and customised the OS according to your needs. Now your friend does not have an Internet connection and needs all these additional packages on his system. Or, all of a sudden, the OS has crashed, and you don't have enough time or resources to repeat all the steps.

Under the circumstances, back up your partitions every time you update your system—this back-up can be shared with any other system or used to restore the OS when it crashes.

Partimage helps you with these kinds of back-ups. It can be installed on many distros via package managers or, preferably, you can run it off a live CD. It is a part of the System Rescue CD.

Here's how you back up a partition:

1. Run the **partimage** command from a terminal.
2. Select partitions for which a back-up needs to be taken.
3. Specify the image name along with a suitable path. Another partition/removable disk has to be mounted in order to store the image in it. (In case of single partitions, or if none of the additional partitions have enough free space, you can use NFS- and Samba-shared mounts to save images on a remote computer.)
4. Opt for 'Save partition...' and go to the next screen (short cut: F5).
5. In the next screen, select the compression level, split mode and finish with other default options (Figure 6). Split mode (auto) means that when a volume in which the image is being stored is full, the image is split into multiple parts. Otherwise, fixed size parts can also be created for the convenience of storing on CDs/DVDs.
6. You'll be prompted to give a description of the image for later use, and one more confirmation screen before copying starts.
7. The copying process will start now. Options are available to cancel, pause and resume the operation.
8. A confirmation screen appears after the completion of the process.

The following are the steps for restoring a partition:

1. Run **partimage** and choose the partition that needs to be restored.
2. Give a path to the image (from a mounted partition/removable disk).
3. Opt for 'Restore Partition...' and go to the next screen (shortcut: F5).
4. Opt for defaults and hit 'Continue' (F5).
5. Now, the image description should appear, which is given at the time of backup, and one more confirmation screen before restore starts. Figure 7 shows a typical restoration in progress.
6. A confirmation screen appears after completion of the process.



Note: When restoring an image, the size of the partition that's being restored should be greater or equal to the size of the original partition on which the backup is taken. However, the target partition is accessible up to the size of the original partition only, in case it is larger. Even though the image size is smaller with only used blocks, the target partition can't be smaller than the original one.

```

Partition to restore:...../dev/sda1
Size of partition to restore:.....20.00 GiB = 21476173824 bytes
Current image file:...../mnt/sda5/image1.tar.gz.000
File system:.....reiserfs-3.6

Partition was on device:...../dev/sda1
Image created on:.....Wed Aug 12 02:33:28 2009
Size of the original partition:.....20.00 GiB = 21476173824 bytes

Time elapsed:.....5m00s
Estimated time remaining:.....38m00s
Speed:.....45.00 MiB/min
Data copied:.....3.25 MiB / 32.70 MiB

```

Figure 7: Restore in progress



Tip: When an image is created for a partition holding the MBR, only MBR can be restored to the system from that image at a later time, by choosing the option “Restore an MBR...” from the initial screen (refer to Figure 5).

Some machines may come with a single partition or not enough space on any of the existing partitions. One solution that was mentioned earlier is to use NFS/Samba to store images on a remote machine. Another simple solution is to use the *partimage-server*. A remote machine running *partimage-server* can be used to store images whose IP address and port can be specified during backup/restore of a partition.

Install *partimage-server* on any distribution and run it using the command *partimaged*. By default, images are stored in */var/lib/partimaged/*. You can use the *-d* or *--dest* options to specify any other directory to store images on the server. Similarly, you can change the default port number 4025 with the *-p* or *--port* option.

In order to connect to *partimage-server* during backup/restore operations, simply opt for ‘Connect to Server’ in the initial screen (refer to Figure 5) and enter the IP address of the server and the encryption mode. Also change the port number to connect to if the server is using a port other than the default one.

System Rescue CD

A good utility that combines many of the tools discussed here in one CD is the System Rescue CD distro. It comes with many additional utilities like:

- An SSH server to allow other systems to log in remotely and copy files
- Filesystem checkers, scanners, mount and view contents of other partitions
- *mkfs* to create a new filesystem and *sfdisk* to set up partitions
- *parted* and *GParted* for partition manipulation
- Firefox/Dillo for instant browsing
- PXE server capability

Tips and tricks

1. To set up a network manually on any live CD, use the following code:

```
ifconfig eth0 192.168.1.1
```

```

Starting restore scripts...

***** SystemRescue-CD ***** 1.6.0 ***** tttyL6 **
http://www.sysresccd.org/


• Type net-setup eth0 to specify ethernet configuration.
• If you have a corporate ethernet network, you can configure by hand:
  - ifconfig eth0 192.168.x.x (your static IP address)
  - route add default gw 192.168.x.b (IP address of the gateway)
• To start an ssh server on this system, type /etc/init.d/ssh start.
• You will need to create an user or to change the root password with passwd.
• Available console text editors : nano, vim, gemacs, joe.
• Graphical web browser is the console: links -g www.web-site.org.
• WARNING : Never mount anything on /mnt! It would freeze the system.
  Use mkdir /mnt/mkdir and mount on /mnt/mkdir instead.
• Ntfs-3g : If you need a full Read-Write NTFS access, use Ntfs-3g.
  Mount the disk: ntfs-3g /dev/sda1 /mnt/windows
• Graphical environment : use either Xorg or Xvesa.
  Type wizard to run the graphical environment (or startx but it may fail)
  X.Org comes with Window-Manager and you can use several graphical tools:
  - Partition manager : gparted
  - Web browsers : .....firefox-2.0 and dillo
  - Text editors : .....gvim and leafpad
02:50 root@sysresccd /root %

```

Figure 8: System Rescue CD

```
ifconfig eth0 up
```

Use a different device number other than *eth0*, if required. On the System Rescue CD you can use the *net-setup eth0* command instead.

2. While running *g4u*, make sure of your disk IDs by running the *parts* command first. For example, if the original disk is detected as *wd1*, and the second one as *wd0*, running *copydisk wd0 wd1* will mean losing all your data in the original disk. Or a safer method is erasing all partitions in the target disk before cloning, by using *GParted* or any similar utility. Now *parts* on the target disk will not show any partitions, whereas on the original disk, partitions will be listed.
3. In the System Rescue CD, if your *partimage* is already running, go to another virtual console (using *Ctrl+Alt+F2*, *Ctrl+Alt+F3*, etc) for mounting, checking image paths, etc. Use *Ctrl+Alt+F1* to come back to the first console.
4. Don't try to save an image in the same partition you're backing up.
5. As Linux distros are compatible with many configurations, cloned disks or partitions work fine in systems with different configurations also— with the exception of one or two features. For example, an image taken for a Linux distro on a Core 2 Duo system works well with PIII or AMD machines also, only requiring you to tweak */etc/X11/xorg.conf* according to new display settings, if you have a problem with starting the GUI. **END** 

References

- GParted: *GParted.sourceforge.net*
- *g4u*: *www.feyrer.de/g4u*
- Partimage: *partimage.org*
- System Rescue CD: *www.sysresccd.org*

By: Rajesh Sola

The author is a faculty member of the computer science department at NBKRIST, Vidyanagar. He is a FOSS enthusiast and contributor. He has presented talks on VBA interoperability and programming OOo with macros and extensions in many FOSS conferences. You can reach him at *rajesh at lisor dot org*.

Use a GPRS Handset as Your PC's Gateway to the Web



Yes, it's all so easy with a tool called KPPP.


When we buy a GPRS-enabled handset, it comes with a CD that contains software that only supports Windows (or if you're lucky, a Mac). This software helps your PC to use your GPRS-

enabled handset as its gateway to the Internet. Does that mean we Linux users are out of luck here?

Not at all! Linux comes with KPPP, a very handy dialler application for CDMA and GSM phones. If you don't have it, it's certainly available in your distro's software repository.


All you need is the access point name (APN) of your network provider that you anyway use when you get online from Windows. Now let us look at the procedure with a series of snapshots:

1



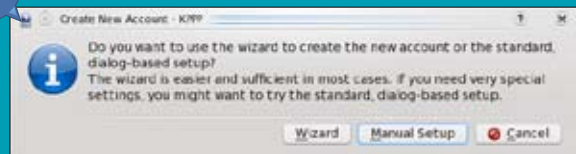
1. Open KPPP. Click on *Configure* to get the *KPPP Configuration* window.

2




2. Click *New*.

3



3. Click *Manual Setup*.

4



4. In the *New Account* window, enter your connection name, i.e., the name of your service provider, like BSNL, Reliance, Airtel, etc.


Click on *Add* and enter your phone number in the small pop-up window.

• GSM user → *99#

• CDMA user → #777

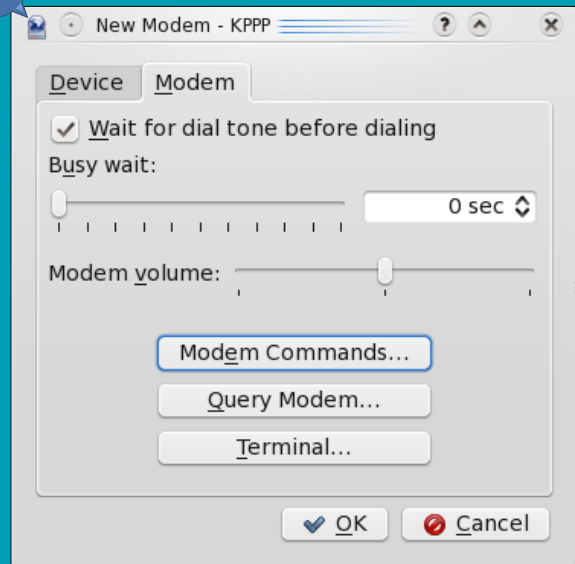
Click the *OK* buttons on the window in which you entered the number and also on the *New Account* window.

5




5. On the KPPP Configuration window, go to the Modem tab and click *New* to get the *New Modem* window.

7



7. Click the *Modem Commands* button.

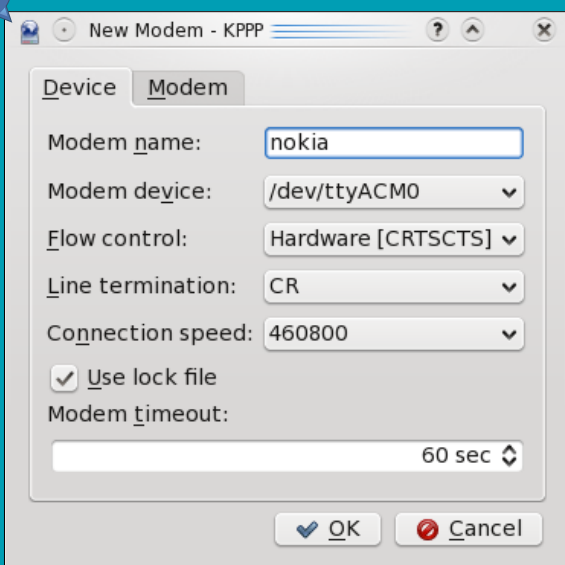
8



8. In the "Initialisation string 2" section, enter: `AT+CGDCONT=1,"IP",accesspoint name`

Note: If you're a CDMA user, leave the "Initialisation string 2" command vacant.

6



6. Type your modem's name here (say, the name of your handset manufacturer) and select `/dev/ttyACM0` from the drop-down menu beside the modem device option. Now, go to the *Modem* tab.


9

Click *OK* (on the three windows) till you get back to the primary KPPP screen (shown in step 1) where you'll enter the login ID and password.

- GSM users enter the number zero (0) for both the login ID and password.
- CDMA users need to enter their mobile phone number in both the login ID and password fields.

10

All you need to do now is click the *Connect* button. Wait for 2-3 seconds, and you will be connected to the Internet!

So, there you go. Was it so hard to do without the software provided by your handset manufacturer? **END** 

Dr. CNS Murthy & Ambuj Dubey

Dr. Murthy is the director of Chamelidevi Institute of Technology & Management (CITM), Indore. He has been working on Linux for the last 20 years. Ambuj is a final year student pursuing BE in Electronics from CITM and loves to play with Linux distros and FOSS.

Video Editing on GNU/Linux

It's EASY
with Kdenlive



Getting started with video editing could be easy with this non-linear video editor called Kdenlive. Here's how.

I wrote a quickie on Kdenlive last time and wanted to do a more thorough job this month because this is one killer video editor. If you like to make movies and show them to your friends, or even need software for a presentation, Kdenlive is your candidate. If you want to create something that looks professional, I can tell you that while it may not stand as tall as Final Cut Pro, Kdenlive comes quite close to Sony Vegas Pro or Adobe Premiere.

It doesn't matter if you have a Web cam, a mobile phone camera or a high-definition camera; Kdenlive is ready with all the punches. The software supports so many formats that it's not possible for me to give the list without missing out something (for the full list, visit kdenlive.org/about-kdenlive/audio-and-video-formats).

Window shopping

Once you install Kdenlive, the rest is as easy as a walk in the park. In Figure 1, you can see multiple panes. The first window is the project tree that shows you the clips/images/audio files you are using in this project. You can just browse for the files here and bring them to the timeline, or a simple drag-and-drop will work too.

Remember one thing: when you start a project, do not change the location of the files being used. Otherwise, Kdenlive will not be able to find those files, which will result in a broken project. Of course, you can relocate the files using a file browser. But it is better to create folders to save such clips. That way, in case you have to move the whole project from one machine to another, you can do so easily.

Anyway, first, you have to start a new project. You can create one by accessing *File->New*. Make sure you create it in a master folder as it will create numerous sub-folders, which it will use for editing purposes. Now, save the project in the same folder. Remember, this is not the movie, but only the project file that keeps all the references. Here you can select the video profile (Figure 2) and the format in which you want to create and edit your movie. You can also import clips directly from your video camera. However, the moment you unplug the camera, the references to the clips are gone. So, either keep the camera attached or save those clips locally, beforehand.

The second window is the 'Effects window' where you can use plug-ins to add visual effects to your clip. The third is the 'Clip monitor' where you can preview the working




Figure 1: Kdenlive comes with five panes by default

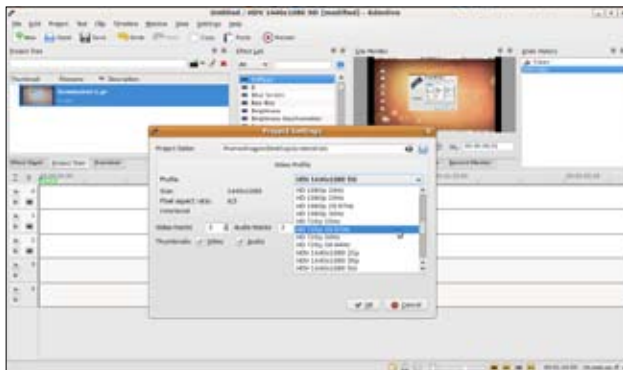


Figure 2: Selecting the video profile

projects and clips. The fourth window is 'Undo history' where you can see all the latest actions made by you, and you can also undo them. The last and the most important window is the timeline. This panel is what makes movies.

The timeline

You can drag your clips to the timeline. Here, you can right click and view all the options for what you can do with the clips (Figure 3). I would suggest you go for 'Split Audio' so that you have audio and video on separate tracks, which gives you more control over the editing procedure. You might wonder about the need for this. Well, let's assume you have taken a shot of a conference and you now want to add your commentary to it. You will need to cut out the live audio from the conference and add your layer of commentary. Moreover, you can also boost the audio if you want to show some response from the audience. It is always better to have independent control on your media.

Here, on the timeline, you can add as many audio and video tracks as you wish. Now, in case of the footage shot at an event, you might wish to make a documentary. You could also take pictures of interviews elsewhere, or of different stalls and activities. To put them together, avoid the same track -- this causes confusion and overlap. Always use different tracks for different locations/activities and then you can place them on the timeline according to your script.

Keep one track for the V/O (voice over) or your commentary and another track for background music. This way, you will have full control over your project and you will




Figure 3: A right click gives you all the options for your current timeline

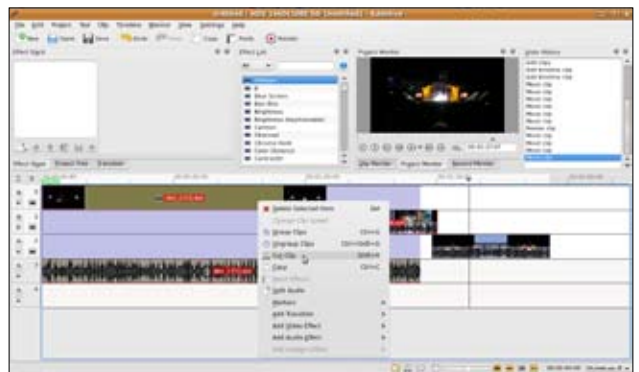


Figure 4: Cut and paste on the timeline

be able to edit shots easily without disturbing other clips. By right clicking a clip, you can also cut it and insert another clip (Figure 4), on a different track. You can also right click to add a 'Guide', which will be a horizontal line so that if your other clips are somewhere down the timeline, you can place them at the correct point. The beauty of Kdenlive is that, just like any other non-linear video editor, you can add, delete and insert clips anywhere on the whole timeline.

You can also select video transition effects from the second video to apply the right transitions. It could fade from black (good for the opening shot—see Figure 5) or you could also create 'title clips' by going to *Project->Add Title Clip*. Here, you can type the text, add the background, colour it, or do what ever you want. Once the clip is ready, save it and it will be visible in the 'Project Tree'. Just drag the clip to the desired place. You may also control the duration of the clip by simply stretching it.

Rendering

Now, all you have to do is render the video in your desired format (Figure 6). But here is the tricky part. There could be a lot of unused clips and audio files on your timeline and if you render this project, all that unwanted material will also get added to your film. So either remove those clips from the timeline or use the 'Lock Track' option to lock the tracks where they are on the timeline, so that they don't get moved while you are fiddling with other tracks. This is very important while rendering a project.

To render your project, look out for the green bar at

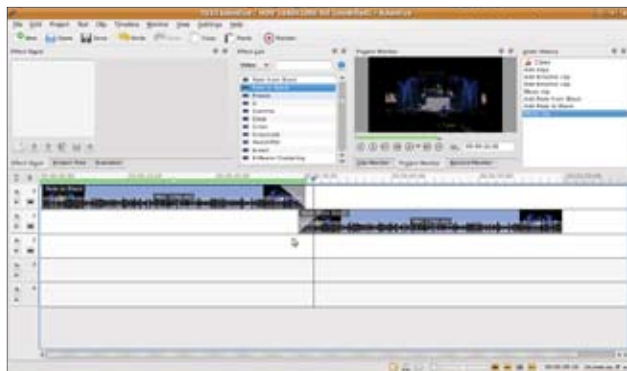


Figure 5: Creating the 'fade from black' video transition effect

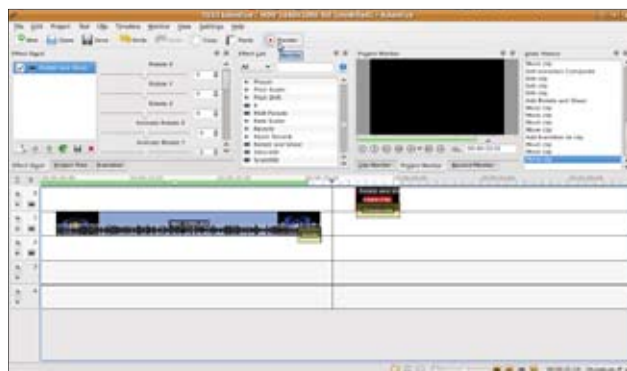


Figure 6: Rendering the video




Figure 7: Various render options

the top of the timeline and stretch it from one end to the other end of your project. Rendering will export only this portion of your timeline. Now, when you click on the **Render** button, you will see the **Options** window (Figure 7). Here, you will have to select the format in which you want to export or render your movie. This selection will recommend the correct format type for the medium in which you want to burn your film.

Then select the location that you want to save the file in, under the 'Output File' option. Make sure you have enough space on that location as exported/rendered media can be huge, depending upon the format you have chosen. Then select the 'Selected Zone' instead of 'Full Project', unless you want the former, and click on 'Render To File'.

There! You have exported your very own video using Kdenlive. I am impressed with Kdenlive because unless you push it to the limits by adding a lot of clips while moving them around, it will not crash. I have even tried this for two days and it crashed only twice, and only when I dragged a lot of clips on the timeline.

Some film making terminology

- **Cut-to-Cut** – When you add one clip after another, without using any effect
- **Fade to black or fade from black**
- **Fade in and fade out** – Audio volume increasing from 'none' and audio decreasing to 'none'
- **Edit Point** – Whenever you are shooting a video, always start shooting at least 5-10 seconds before the event/activity starts and shoot at least 5-10 seconds extra after the activities end. This way, when you go to your Kdenlive editing table, you won't bang your head that you cannot find the edit point. The edit point is that specific frame that goes well with the last frame of the previous clip, or the first frame of the next clip.

Bottom Tool Bar


- **Selection tool** for normal selection
- **Razor tool** – select this to cut a clip
- **Spacer tool** – this can move all the clips together on to the timeline.



Tips:

1. Never shoot against the sun, it may damage the CCD of your camera.
2. Always take care of the White Balance. If you are shooting in different lights, please adjust the White Balance of your camera as every light has its own colour temperature. If you do not do this while shooting, you might find that the colours of the object have changed in the movie.
3. You can read more about the grammar of film-making here... <http://www.aber.ac.uk/media/Documents/short/gramtv.html>

So, Kdenlive is an awesome and pretty powerful tool for film lovers. If you like making films and think FCP or Sony Vegas Pro are better, instead of wasting \$1000 on buying them, get together some friends who need these tools and collect the money they plan to spend, from them. Thus, if there are 10 friends, you'll have \$10,000 and you can use that to fund Kdenlive's development instead!

You can find the video I created using Kdenlive at <http://www.youtube.com/watch?v=4uHZipQwyPA> **END** 

By Swapnil Bhartiya

A Free Software fund-a-mental-ist and Charles Bukowski fan, Swapnil also writes fiction and tries to find cracks in a proprietary company's 'paper armours'. He is a big movie buff, and prefers listening to music at such loud volumes that he's gone partially deaf when it comes to identifying anything positive about proprietary companies. Oh, and he is also the assistant editor of *EFYTimes.com*.



Addressing the Linux Talent Shortfall

Over the last few years, Linux has proved its worth in every vertical imaginable. Yet, corporate India is caught in a bind with Linux-conversant talent being virtually unavailable. Shenoy Systems has stepped in to speed up the process of getting talent industry-ready, especially in the embedded systems arena.

Linux is being hailed as the long-awaited magical wand for corporates looking at lowering the cost of building products.

But the magicians seem missing, with not many developers skilled to employ Linux in building devices. Naveen Shenoy, MD and founder of Shenoy Systems, hopes his training solutions on Linux and embedded systems can fill the gap between corporates' demand for Linux expertise and the talent pool available, making the Linux magic happen!

First dip: embedded system solution provider

However, finding that perfect training spell wasn't instantaneous for Shenoy Systems. It involved a year-long journey, from an embedded system service provider to training solution provider, with one common thread—addressing the immediate need of the embedded industry. Recollecting the initial plunge, Shenoy says: "When I was to start Shenoy Systems, I was looking at addressing an immediate need of the embedded devices sector. With 11 years of experience in embedded systems, I felt there was a demand for an exclusive embedded system service provider, as companies that offered embedded services

made the service team work on a project-to-project basis. Such a team lacked in-depth knowledge on embedded systems."

This resulted in the birth of Shenoy Systems as an affordable and high-quality embedded systems service provider in October 2007, which offered end-to-end solutions for real-time systems (RTS) and real-time operating systems (RTOS), design and development of firmware hardware abstraction layers and device drivers. The earliest clientele for their consultation/services around embedded systems included L&T, Tata Elxsi, Aspect Technologies and D-Link.

However, the initial journey wasn't free of obstacles. The company faced funding problems and hitches in the timely delivery of solutions. "Timely delivery of solutions was a problem mainly because Shenoy Systems was a one-man army then. I found it difficult to find appropriate talent to support my endeavours," says Shenoy, who has technical experience with Wind River, TATA Elxsi, and Park Controls.

Training solutions: The plunge again

But the obstacles turned to be a blessing in disguise. The difficulty faced by Shenoy to find well-equipped talent made him realise the need for training solutions in the

embedded systems industry, especially with the rise of embedded Linux. This realisation led to the germination of the idea of providing Linux-based training solutions. Personal experience further strengthened Shenoy's belief in the need for such a solution. Shenoy recollects: "When I started working on Linux six to seven years back, I had a tough time as I learnt everything myself and then used Linux to build embedded devices. The process was time consuming; it took double time to build the embedded Linux product. Considering time is money in our business, it was a loss for the company," says Shenoy. He adds: "If I had adequate training I would have been able to ensure timely delivery of products employing Linux."

The lucrative business model around training solutions was an added attraction for Shenoy. "We decided to take the plunge in training solutions as initial costs were low, plus the business looked promising. We have no regrets as not only have we broken even but also earned profits in a short span," says Shenoy.

A one-stop-shop for embedded Linux training

In April 2009, Shenoy Systems began offering training solutions in generic Linux, OS internals, BSP (Board Support Package), RTOS and device drivers for individuals, corporates and educational institutions. However, bagging the first orders for training solutions didn't come easy for Shenoy. They had to prove their worth when compared to academicians and well-established Linux solution providers.

The company promoted the practical experience of its team as its USP to sell training solutions. "We at Shenoy weren't academicians with only theoretical knowledge. We had the industry experience by which we could enable professionals to use Linux for building products. We could even provide hands-on experience. This became our USP and helped us crack our first deals," says Shenoy.

Further, the price competency of the training solutions attracted corporates and developers. "While the cost of Linux training courses for companies was over Rs 60,000, we charged one-third of it. Similarly, for individuals or fresh engineers/students we offered training at a lower price," says Shenoy. Finally, after a month-long struggle, Shenoy Systems got its first clients for training solutions. "Our first clients were L&T, EmSys and Tata Elxsi where we held training in embedded

systems, RTOS, embedded C, embedded Linux and generic open source courses," says Shenoy.

Shenoy today: always innovating

Today, Shenoy Systems has an impressive list of corporate clients that include Arrow, HCL, D-Link, HP, Aspect, Samsung, Tata Elxsi, and L&T Infotech. Further, they have two partners—AayuSmart and Oasis Technology. Currently, the company is involved in a medical device project. "We are working with AayuSmart, our partner, on building medical devices like glucose devices, BP apparatus, ECG machines, etc. These devices will make patient reports accessible to a doctor at anytime, from anywhere. The small 4x6 inch (10.2x15.2 cm) ECG machine can be connected to a cell phone and data can be uploaded on the Internet through GPRS, SMS or Bluetooth. This will be a step ahead from the telemedicine that makes data on the patient available within the hospital premises," says Shenoy.

The company is constantly working on improving the delivery of existing solutions. "Apart from classroom-based training, we impart audio-ware training that provides a student flexibility to come and learn as per his convenience. Every candidate gets a DVD player and notes. The audio tracks have detailed explanations on the topic, while the notes carry the diagrams and theory aspects. In case of doubts, the on-premise trainer clarifies them instantly," says Shenoy. Such innovations ensure quality learning, and break time barriers. They are extremely beneficial for working professionals looking for a career

shift to embedded systems, embedded professionals trying to gain expertise in the domain, and fresh engineering graduates or diploma holders.

Elaborating on the merits of the audio-ware programme, Shenoy says: "We also provide hands-on experience that enables candidates to walk the talk with concepts they learn from audio. The audio courses are currently available for embedded systems, RTOS and soon shall be started for Linux," says Shenoy.

However, maintaining quality doesn't come easy for Shenoy. Security of information/ training material is crucial. "There is no scope for tampering with the training recordings. The DVD player is sealed with a sticker and our staff ensures everything runs smoothly. We are also working on disabling USB ports so that the student concentrates only on the lessons," says Shenoy.



"With an improvement in support, services and training around Linux, we predict exponential growth for Linux-based devices."

—Naveen Shenoy, MD and founder of Shenoy Systems

Factsheet

Established in	October 2007
Founder	Naveen Shenoy A
Partners	Aayu Smart, Oasis Technologies
No of employees	5 fulltime professionals and 6 retainers
Website	www.ShenoySystems.com
Address	#3/2, 2nd Floor, 9th Main, Banashankari 2nd Stage, Bangalore-560070
Email	info@ShenoySystems.com
Phone	9481471581 (Mobile), 080-26716446 (Landline)

Side-stepping pebbles en route

"Mindset remains a problem. For instance, if a vendor proposes Linux to a customer, the customer will check with their administration team, IT suppliers etc, for whom Linux is not the default OS. Thus, it becomes a very tough battle to be fought and won," says Shenoy, who has now deployed the strategy of collaborating with educational institutes to offer Linux training at the root level. They conduct seminars at colleges and talk to lecturers, helping the latter understand industry requirements.

Shenoy Systems still faces reluctance among freshers and experienced people to join them. "Freshers want fatter pay packets while experienced people are hesitant to take on the numerous challenges a small firm faces," says Shenoy.


Betting on embedded Linux

Currently, Shenoy Systems has two vacancies for handling their audio training module, handling projects, etc. And what are the skillsets required? "We are looking for a person with electronics or electrical background for this profile," says Shenoy.

In the future, Shenoy anticipates a requirement for about 15 engineers. He feels embedded Linux holds tremendous opportunities for developers, as by 2011 the

embedded device market will see exponential growth with most devices running on the Linux OS. However, "Training is a must for developers. It helps you learn from others' mistakes," advises Shenoy.

With the mobile, automotive and medical space increasingly using embedded Linux, it is a promising area for entrepreneurs too. "There is a need for Linux embedded devices to be developed in India. We expect a huge demand for such devices from the electronics, networking and medical sectors," says Shenoy. "An embedded systems startup should work on getting partners and developing an entire ecosystem around embedded systems. Only then will the solutions click with the customers."

The company is optimistic about the demand for Linux in the future, with the convergence of technology and the obvious cost benefits of using Linux in embedded devices. "The cost advantage Linux provides in making devices will continue to attract people to Linux. With an improvement in support, services and training around Linux, we predict exponential growth for Linux-based devices," winds up Shenoy on a bullish note. **END** 

By: Vanisha Joseph

The author loves to experiment and writing for LINUX For You is her latest experiment. So, beware! Just a minute, she also happens to be a journalist during the day.



redhat®

TRAINING PARTNER

RHCE / RHCSS Exam Centre

Hundreds of hardcore Industry Professionals Vouch for it.



RHCE only @

ADVANTAGE PRO

- Pioneer in LINUX training since 1999.
- Chennai's #1 RHCE training Centre
- 100% results every month.
- Many centum students.
- Also get trained on My SQL - CMDDBA, CMDEV, PHP, Perl, Ruby, Python, Ajax...

Next RHCE/RHCSS Exam
 14, 15 & 30th September'09

Don't wait.
Register Now.

Let the success vouch for you!



ADVANTAGE PRO

IT TRAINING DIVISION OF
Vectra Technosoft Pvt. Ltd.
Open Source Academy

Regd. Off: Wing 1 & 2, IV Floor,
Jhaver Plaza, 1A, N.H. Road,
Nungambakkam, Chennai - 34.

Ph : 28263529 / 3530 / 3540

Telefax : 28263527

E-mail : enquiry@vectratech.in

www.vectratech.in

Develop a Simple Download Scheduler



Here's how to write a few simple scripts and then schedule them to run at a specified time using the *at* utility.

Many of us have an Internet connection that offers unlimited downloads at odd hours, usually around midnight—mine is between 2 a.m. and 6 a.m. What most of us do is wake up late at night to download whatever heavy files we wish to, in order to capitalise on this unlimited bandwidth. However, it becomes quite inconvenient to get up every day in the wee hours and set up things in order to download. Then there's the additional headache of again getting up when the time slot for unlimited bandwidth ends and cancel the downloads. If only there was a way to do these things automatically without the need to manually start and stop the download processes.

In this article, we'll talk about how we can write a shell script to create a simple download scheduler for domestic purposes. Additionally, we'll explore a way in which we could make the PC automatically connect to the Internet, provided we schedule the

job at the required time (scheduling will be described later).

Note that by writing a script, we'd in no way write any functions and programs, but combine some features of the following utilities:

- *lynx-cur* – a command-line Web browser
- *wget* – the non-interactive network downloader
- Bash shell – for running our script, of course
- *at* – a utility used for scheduling jobs

Make sure your PC has all the above utilities before going further. *wget*, *at* and Bash shell, might be there already. However, the same might not be the case with *lynx-cur*. If it is not present, run the following command:

```
sudo apt-get install lynx-cur
```

How to go about it

The following tasks will be performed by our scheduler—all automatically:

1. Connect to the Internet at the scheduled time
2. Download the file present at the provided URL
3. Log out from the Internet service at the scheduled time
4. Shut down our system at the scheduled time

Before we dive into the subject, there are two things to make a note of: make sure you log in to the system as the root user, and save all scripts and text files that we will create in a single directory.

Also, the first and third tasks listed above—connecting to the Internet and performing the logout—apply only for those connections that demand users to log in by submitting their user name and password into a Web page located at the ISP's server, and to log out by clicking on a logout link (for example, Sify). People with always-on connections (like Airtel, BSNL, etc) need not worry about these two tasks as most of these connections ensure connectivity as soon as the modem/router is connected to the system. Those belonging to this category only need to worry about downloading and shutting down the system at the right time.

We will run three independent scripts: *start.sh*, *logout.sh* and *shutdown.sh* — each from three different terminal windows. *start.sh* will control logging in to the Net and starting the download. *logout.sh* will log out the system, and *shutdown.sh* will shut down the system.

A note on lynx-cur

lynx-cur is a command-line browser—there's no mouse and no graphics here; you can only browse using keystrokes. If we are able to create a file that contains the sequence of key-strokes required to navigate a particular Web page, and if we are able to effectively transfer this data to the browser, then logging in and out are easy.

This is where *lynx-cur* comes to our aid. It has the striking feature of reading the key strokes from a file and then sequentially applying them on the given Web page, creating the same effect as if we were manually doing it. All you need to do is to provide an option: *-cmd_script=<file name>* while invoking the browser to browse a page. Lynx will read the key strokes from the specified file. For example:

```
lynx -cmd_script=<file name> <URL of the web page>
```

So the problem of an automatic login and logout is resolved provided we have a file that has the key-strokes data logged exactly in the same way the browser would like to read and interpret it.

But how do we create such a file? Lynx has an answer for this one too. Just like the *-cmd_script* option, Lynx provides another option called *-cmd_log*. With this option, the browser will log the key strokes we execute while browsing to a specified file (after it gets invoked and till we quit). For example:

```
lynx -cmd_log=<file name> <URL of web page>
```

Logging in

Command-line browsing in Lynx is all about moving the cursor in between hyperlinks and text fields (user name and password) in the login page by using the *Tab* key, and entering text into the text fields. In order to follow a link, we select the link by using *Tab* and press *Enter*. So it is understood that we will keep a text file created by Lynx, which it will use while running the *start.sh* script. To create the file, run the following command:

```
lynx -cmd_log=dir_in <URL of your login page>
```

The browser will open the login page. Now move the cursor to the text field where we are supposed to enter the user name and password, and enter the necessary data. After that, try to log in by moving to the login link and pressing *Enter*.

Remember that the process of logging in may vary with ISPs and the type of connection you have. For example, those with always-on (aided by a dedicated router/modem) connections can skip building the *start.sh* script, as they will get automatically logged in after they boot their computer. Here I am sticking to the core idea of a Web-based login method—i.e., reading keystrokes from a file, and even that, only if users have to submit their profile to authenticate their identity. Note that the name of the file given here is *dir_in* keeping in mind it contains directions for the browser to log in.

Now that you have logged in, you can close your browser. Press *q* and it will ask for confirmation. Enter *y*. That's it. Our *dir_in* file is ready.

You can see that *dir_in* has got all the keystrokes you've made. Also, you can see your user name and password. So, make sure you have changed the read, write and execution permissions of the file:

```
chmod 700 dir_in
```

Hereafter, if you need to log in, you can simply run the following command:

```
lynx -cmd_script=dir_in <URL of your login page>
```

...which will fill your authentication details, submit the data, and will finally quit from the browser to return to the terminal.

Finally, save the above command in a file named *login.sh*.

What if you're not able to connect the first time you invoked *login.sh*. In that case, you will have to run *login.sh* repeatedly till the connection is established. For that, embed the above command in the script within a *while* loop as follows:

```
#login.sh
#!/bin/bash
status=1
while [ $status -ne 0 ]
do
```



```
lynx -cmd_script=dir_in <URL of login page>
status=$?      # $? returns 0 if previous command was executed
successfully.
done
```

That makes our final *login.sh* script file ready for use.

Downloading

Now that the system is connected, the next step is downloading the required files. We will use *wget* as our download utility, which accepts the URL of the file to be downloaded as its argument. We can invoke *wget* as follows:

```
wget <URL>
```

You can specify the folder to which *wget* saves the downloaded material. Simply add an option, *—directory-prefix=<target folder>*. For example,

```
wget —directory-prefix=/media/new_volume <URL>
```

Here, */media/new_volume* is the mount point for my backup drive, which contains all my downloaded files.

However, before we create the download scheduler script, we should create a collection of URLs of the files we wish to download, one after the other.

Let's now write another script called *download.sh*, which contains a set of *wget* commands pointing to the required URLs. Here's an example:

```
#download.sh
#!/bin/bash
wget -directory-prefix=<target folder> <URL of file 1>
wget -directory-prefix=<target folder> <URL of file 2>
##[append as many wget commands as you wish]
wget -directory-prefix=<target folder> <URL of file n>
```

The above script is, naturally, the most important part of our *start.sh* script.

It includes *login.sh* and *download.sh* and, of course, is the script you need to run in any one of the three terminal windows as mentioned earlier. The function of *start.sh* is simple—run *login.sh* and then *download.sh*. Here's what it should look like:

```
#start.sh
#!/bin/bash
bash login.sh
bash download.sh
```

Logging out

After downloading the necessary files, the next job is to log out. So we are on our way to write the *logout.sh* script. The steps to log out are similar to logging in. We will create a file called *dir_out* to store the keystrokes we need in order to log out.

As with logging in, the procedure to log out may vary with ISPs and connection-plans. For instance, in my case, I need to re-submit my user name and password in order to log out. My ISP doesn't have an independent logout page. So when I open my browser with the URL of my ISP, even when I am online, I'll have to enter the user credentials to access the logout page. However, for others, this might not be the case. If your ISP's servers can auto-detect the connection status and pull up the logout page instead, when you invoke the browser with the ISP's URL it's easy to create *dir_out*, as you only need to run the following:

```
lynx -cmd_log=dir_out <URL>
```

...and keep hitting the *Tab* key till you reach the 'logout' link; finally, press *Enter* to log out.

Those in a similar situation can run the same command as above, resubmit a login ID after filling the text fields, and submit the logout after highlighting the logout link by moving over preceding links in the page. Now quit from the browser and provide the confirmation. Our *dir_out* file has now been created.

Again, it is possible that the logout isn't successful at the first attempt. So let's bring in the magic of the *while* loop again and our final *logout.sh* script should look like what follows:

```
#logout.sh
#!/bin/bash
value=1
while [ $value -ne 0 ]
do
    lynx -cmd_script=dir_out <URL of ISP>
    value=$?
done
```

Shutting down the system

After essential downloads and a successful logout, it would be wise to shutdown the system—we should save electricity, seriously. Here, we can make use of the *halt* command. Note that you can't be careless while you invoke *halt*. It will initiate a forced shutdown and all your unsaved data will be lost. Also, it can only be executed by the root user. Save the command in a separate file called *shutdown.sh* with the following data:

```
#shutdown.sh
#!/bin/bash
halt
```

That's it; all the essential script files are now ready. It's time to put these in a schedule.

Scheduling jobs

So far, nothing has been discussed about how you can run a job (script) at a specified time. This is possible by means

of a utility called *at*. When invoked with the option *-f*, it will run a given script at a given time. The following is a typical way you can use the *at* command:

```
at -f <script name> <time>
```

Here <time> can be specified in different ways, like:

- 11:45 PM
- now
- now+5 minutes

To get more details on *at*, check out the man page.

It is worth mentioning that the output of jobs scheduled using *at* will not be tagged with STDOUT (i.e., the monitor) which means that the process will run only in the background. We will be able to know whether the run was successful or not by checking the mailbox of the person who used *at*, or by using the system monitor utilities like *ps*.



Note: The mailbox here is different. In the current context, it's the e-mail sent by the system to the person (user) who has scheduled the job. If users have installed a utility called *mail*, the output will be written to a text file in the */var/mail* directory. It is not compulsory that you need to have *mail* installed, but for those who are interested in scheduling more jobs using *at*, *mail* is a useful utility.

Let's suppose our schedule looks like the following:

- login and start download at 2.10 a.m.
- logout at 5.45 a.m.
- shutdown at 5.50 a.m.

What we need to do is schedule our scripts (*start.sh*, *logout.sh* and *shutdown.sh*) at the respective time. Open three terminal windows and run the following:

```
# at -f start.sh 2:10 AM ##in the first window
# at -f logout.sh 5:45 AM ##in the second window and
# at -f shutdown.sh 5:50 AM ## in the third window
```

Continuing from stopped downloads

If the download was not over within the time span, a truncated file will be created at the destination folder at the time of logout. This is true for large sized files. On those occasions, you can continue from where you've stopped at a convenient time, by running *wget* with the *-c* option:

```
wget -c <target folder> <URL>
```

If there is a file with the same name as the one specified within the URL in the target directory, *wget* will request the server to continue retrieval from an offset equal to the length of the downloaded one.

Improvisations that we can incorporate

Now, our scheduler can be made to run successfully under normal conditions, which of course includes an


undisrupted connection with the server. For those sites that compulsorily require you to sign in to access the Net, you can simply invoke *lynx-cur* by making it read keystrokes from a file (as we've done earlier) for signing in. In case of disturbances that result in disconnection while downloading, you can ensure a smooth download by making provisions in the script to reconnect. This can be done by writing a function as follows:

```
#download.sh
#!/bin/bash
function do_it {
    wget -t 1 --timeout=60 -c --directory-prefix=/media/new_volume/
yesterday $1
    status=$?
    count=$(( $count + 1 ))
    if [ $status -ne 0 ] && [ $count -lt 5 ]
    then
        bash login.sh
        do_it $1
    fi
    count=0
}

count=0
do_it <URL 1>
.....
do_it <URL n>
```


Here, *do_it* is a function which accepts URLs as arguments which *wget* is then made to point at. The *-t* option specifies number of tries *wget* must attempt to download. Its value is assigned as 1. After the first try, *wget* will quit with a status value. This is examined along with the value of an integer variable *count* to decide whether a re-login has to be made.

Integer variable *count* is specified by us, depending upon the number of times re-login must be allowed. Here, I gave the terminating count to be 5, which means download will continue smoothly by logging in for first five disconnections experienced in course of the download of a file. You can also see the *--timeout* option which can be used to specify the maximum time, in seconds, *wget* must wait for a reply from an idle server. Its default value is 900 seconds. It is set to 60 seconds just to save time.

I guess that's about it, for now. Will catch you later when I chance upon some other useful tip.  **END**

By: Sreekanth Balakrishnan

The author is a third-year B. Tech computer science student from College of Engineering (CET), Trivandrum. He is an active member of fsug-tvm and also the FOSS group at his college. He loves shell- and python-scripting, and is an evergreen fan of KDE.



Tiered Storage Infrastructure

Storage, Server and Application Levels

With the increasing volume and complexity of information that gets churned out these days, the need to optimise storage and retrieval becomes an important aspect to consider. In this article, the author sheds some light on each layer of a tiered storage infrastructure — storage, server and applications.

Today, information is being created faster than ever by businesses across the world. The information growth rate is not uniform across industry segments. In some verticals, growth follows a sedate pace, whereas an information explosion awaits other industries like financial services, health care, etc. In such a complex environment of information management, today's IT infrastructure managers face the challenge of enabling the right amount of infrastructure that can help organisations in leveraging information and creatively managing the data centres optimally. And all this needs to be achieved in the most economical way.

On the one hand is the non-linear

growth of data, regulatory compliance and the decreasing cost of storage, and on the other is the impact of the industry slowdown. All of this has combined to further speed up the process of doing more with less. In order to keep up with these dynamics, IT managers are faced with the challenge of cutting IT budgets while risking the probability of not addressing business needs effectively, both from the resourcing as well as the infrastructure perspectives. Inadequacy in addressing business needs includes exposing business-critical data to risks such as loss of data, theft and downtime of business applications.

This is a pain point for the industry and can be addressed by designing a 'tier-ed infrastructure' model, which is a sub-set

of Information Life Cycle Management (ILM) for data centres aiming to bring in a systematic approach to deploy servers, storage and business applications.

What is 'tier-ing'?

A tier-ed infrastructure brings about a systematic approach of providing the right amount of horsepower to fuel the infrastructure and to handle the business needs in the most optimal way. This is done by leveraging the existing infrastructure to ensure optimal utilisation of the equipment, ensuring the right Quality of Service (QoS), and harnessing the infrastructure and the tools to load-balance the resources, thereby providing high availability.

As shown in Figure 1, a 'tier-ing' exercise involves a holistic view of all business applications as well as data centre inventory. This enables IT managers to formulate the right strategy by creating tiered buckets and fitting the applications, server and storage appropriately. This is done based on their functional and non-functional characteristics and has proved to be an effective means to reduce the overall cost and, at the same time, ensure higher grades of service.

Why is tier-ing important?

Tier-ing ensures optimisation of the data-centre infrastructure to handle business data requirements in the most effective manner, which is essentially the process of considering the functional and non-functional attributes of the different infrastructure components to enable the creation of tiers in IT infrastructure.

The following components form the building blocks of a tiered infrastructure model for organisations:

- Classification of data
- Storage, server and network tier-ing
- Data heat map

One of the key considerations of a tier-ing model is that the service should be provided in tiers (varying the levels of service), which fits the purpose and is not necessarily the best of the breed. Various commercially available Storage Resource Management (SRM) applications may be used to move the data to the appropriate tier of storage, as required.

How does tier-ing benefit data-centre managers?

Tier-ing enables IT managers to understand the requirements of the application data in a holistic way. For example, live production data from a financial services application needs greater bandwidth on a high-performance network with a high Quality of Service (QoS) to be de-staged at the back-end storage, providing equivalent performance levels, whereas data that needs to be archived can have a lower priority and a lesser performance requirement from an availability perspective.

Apart from cost-savings, organisations that use tier-based archiving benefit by:

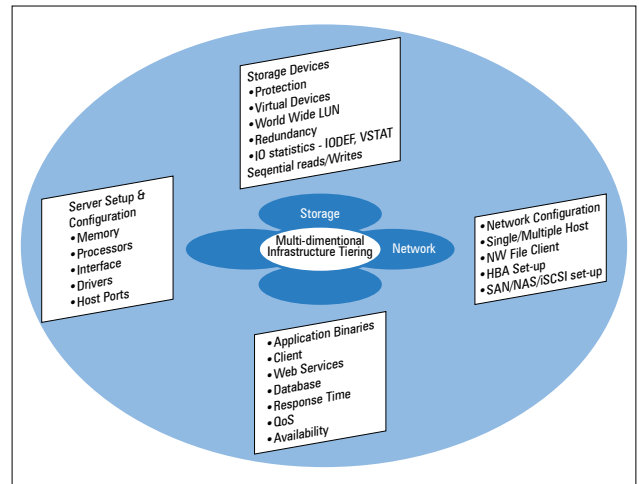


Figure 1: Structure of tier-ing

- Highly reduced turnaround times in the back-up and restoration of their application data
- Application data is optimised for storage
- Availability of guidance on the cost to the application, based on the aforesaid storage tiers
- Dependence of the service levels associated with a given application on the environment in which it is running
- Heterogeneous storage and server management to increase agility, quality of service and the optimising of application performance

On the flip side, there have been several instances where businesses have invested in a tiered storage model, but have failed to reap the benefits and have actually ended up dealing with much larger chunks of unstructured data. This highlights the need for investing in a detailed design and implementation approach to tier-ing, including a thorough benchmarking of the infrastructure.

How to implement tier-ing?

The fundamental reason for the classification of business data in a tier-ed infrastructure model is to address the following industry issues:

- Data classification as one of the important attributes of a comprehensive IT data centre strategy
- The need to reduce IT costs and improve effectiveness
- A means to maximise a business' current infrastructure investments and support its Information Life Cycle Management efforts
- The need to create a data heat map, which comprises the data life cycle management with details on the different phases of data creation, storage, retrieval and archival. The data heat map is also supposed to showcase the changing importance of data to the business at various stages.

The process of implementing tiered infrastructure involves understanding the data requirements of the applications in a holistic way, relative to the performance

requirements of the other applications. One of the main tasks of a tier-ing initiative is to classify the data created based on several attributes such as age of the data, its business value, criticality and its usage.

The value of the information created in businesses changes with time. This is not to mean that it isn't useful. The emergence of regulatory compliance laws such as the Sarbanes Oxley (SOX) and other corporate best practices has necessitated the need for archiving old data from organisational e-mails, sales data, and so on, for a specified period of time. Though this is a low priority segment of data and has a low frequency of usage, it nevertheless needs to be handled in a specified manner, with the right archival policies.

Business-critical data requires a high priority and an equally robust mechanism to ensure the reliability and availability of such data. This type of data includes sales reporting data, IP, clinical research data, and so on, which demand high priority in terms of data back-up, restoration, etc.

Every organisation creates different types of data to support its business. There is primary data, which is critical to the business, whereas the other forms are purely for business support. Some of the secondary and tertiary types of data are created through tools such as BI applications for business analysis. Hence, it is important to categorise the information generated based on the usage of the information for different applications within the organisation.

Case Study - Implementation of tiering

In the case of a particular healthcare company, the data centre was consolidated on a single premium vendor's storage platform. Due to market conditions and increasing pressure to cut IT budgets, the organisation was forced to look at ways to cut excess expenditure on operational costs. This drove the focus towards the business justification of expending premium storage for all applications, irrespective of their requirements.




Figure 2: Storage vs applications

Hence, the management started to take care of new storage requirements on cheaper storage to balance the operational costs on the premier storage.

In spite of the presence of an adequate storage infrastructure, the company faced severe bandwidth issues for about seven of the 45 critical business applications. This was followed by lower utilisation levels on the storage, which was at an unusual level of 45 per cent. The existence of excess capacity as a result of application owners incorporating a 10-30 per cent buffer in their storage requirements over several years, contradicted the fact that the business applications faced availability issues.

This baffled the management and clearly pointed to the fact that there was no clear policy and methodology in place for treating applications' requirements on the merit of their functional and non-functional requirements. The potential danger of risking downtime and data loss on business-critical applications prompted the management to take up a portfolio analysis, with the ultimate goal of creating a data life cycle management model for all the business applications through a tiered infrastructure [refer to Figure 2].

A data heat map traces the various phases of the data's life cycle and the business value that it provides at various stages. This is an important tool to build an effective tiered infrastructure model, as it provides valuable information of the functional and non-functional requirements of the application data at various

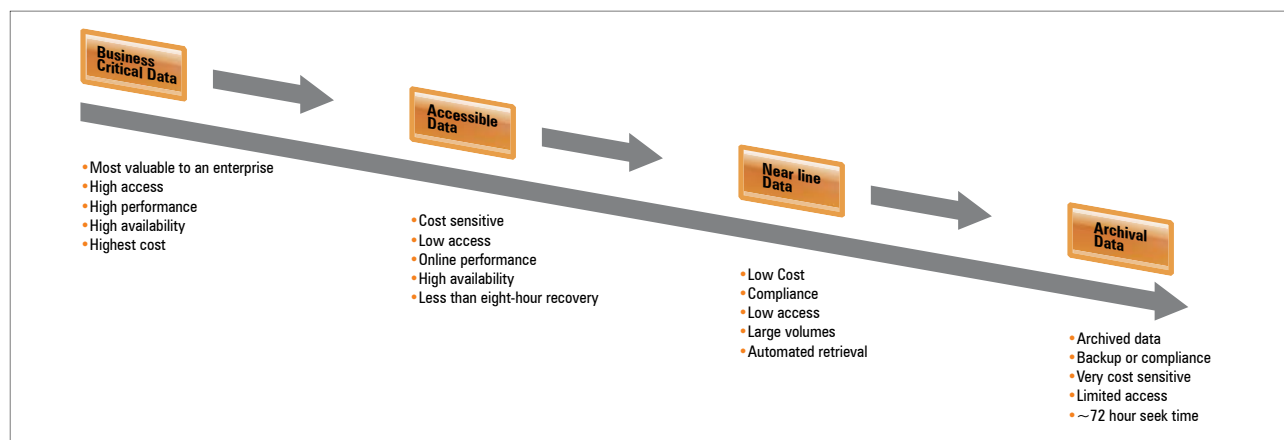


Figure 3: Data life cycle management model for all the business applications through a tiered infrastructure

stages. This can then be aligned with the storage, server and the application infrastructure in the organisation to build the critical parameters governing the definition of the various tiers.


The identification of different categories of data and assigning them to the appropriate types of storage, play an important part in the storage infrastructure tier-ing, which reduces cost as well as optimises the performance levels required by the applications. Functional categories include the type of back-ups, retrieval, type of protection needed, availability requirements, and so on.

Non-functional attributes include performance requirements, frequency of usage of the data stored, and so on. The tier-ing process is ongoing and there are commercially available Storage Resource Management (SRM) applications that can be leveraged to achieve this. In the tier-ing activity, as illustrated in Figure 2, the different vendor storage platforms can be integrated to form a homogeneous pool of storage, using technologies such as virtualisation on top of which a tiered storage model can be implemented based on the criticality of data stored, performance requirements, etc.

For example, Tier 1 comprises business-critical data that has a high frequency of access and needs to be on premium storage that can address the high performance requirements of this type of data. Tier 2 data, which is relatively lower priority, may be stored on media that

guarantee availability, but at a lower Quality of Service in conventional Storage Area Networks (SANs). Tier 3 is low priority data that is not frequently accessed [such as e-mail archives]. This type of data can be stored on cheaper storage where the performance requirements are less.

Key attributes to be considered when building a tier-ing model from a server perspective are the number of processors, availability of memory, number of I/O ports, etc. In a similar manner, application tier-ing can be achieved by separating business applications into multiple tiers based on functional parameters such as application binaries, criticality of data, etc, and non-functional parameters such as performance and the Quality of Service (QOS) needed.

Infrastructure tier-ing is a holistic capacity planning model that aims at assessing the current and future capacity requirements of a business' IT infrastructure and determines a comprehensive tiered model to handle business applications in the most effective way. It is very important to be aware of all the building blocks of the infrastructure when designing a tiered storage model, without which the true benefits cannot be enjoyed. 

By: Venkat Thummi

The author is a senior manager, Technology, at Cognizant.



Calling...

Share your knowledge

Inviting FOSS experts to write articles on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of www.openITis.com (aka linuxforu.com) we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at

lfyedit@efyindia.com

Few topics that top our list:

- ◆ Tips 'n' Tricks for software developers or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools

GIMP for Beginners, Part 3

Toolbox Demystified

The third article in our series on the GIMP will guide you through Toolbox, and help you to work with the various tools featured in it.



The *Toolbox* that appears on the left of the main window is the foundation of the GIMP. Initially, it may appear mind-boggling, especially for a person who has no background or experience with any graphics editing program. While a detailed description of the *Toolbox* is out of the scope of this article, it would be helpful to get acquainted with each tool that comes with it.

Note that this tutorial will begin with a modified version of the *Toolbox* since we had learnt about customising the GIMP interface in the first article under this series. You can get to know more about that in *LFY's* July 09 issue.

Before we get started, we will cover 'selections'. Here, rather than discussing any type of selection tool, we will learn

about the nature of selection.

A 'selection' can be done by using different tools; you can use 'figure selection', 'colour selection' or even 'path selection'. Every 'selection' shows two types of traits, i.e., passive and active.

Passive selection: A 'passive selection' transpires at the time of drawing or editing a selection, and does not allow any drawing or painting while working with it. The passive selection will either vanish or change into active mode if you attempt to draw anything inside it.

Active selection: An 'active selection' is just a normal selection in which you can draw or paint. To distinguish between the two, the GIMP drafts the active selection with continuous blinking/moving dots. Refer to Figure 1 for more accurate details.




Figure 1: Types of selection

The modified *Toolbox*, shown in Figure 2, comes with 33 tools and two colour swatches. Some of these tools have analogous tasks but they tend to achieve them in different ways.

I have segregated the tools into different categories to help you understand them easily. In this tutorial, we will learn the basic aspects and use of the various tools, and will continue learning about them in more detail in upcoming tutorials (refer to Figure 2).

Drop area

The *Toolbox* starts with a faint 'wilber' logo from the top. This is a drop down area, where you can just drag and drop an image and start editing. It's very handy when working with congested desktops.

Selection tools

The true test of any image-editing program is a comprehensive set of selection tools, which have a wide range of features that go far beyond their stated purpose. You can create stupendous effects only with selection tools. The GIMP comes with quite a few selection tools for various purposes.

- **Rectangle Selection:** This is the first selection tool that you see in the GIMP's *Toolbox*. You can vary its settings to create assorted polygons. You can also use this tool to create rectangles with round edges that come very handy while creating buttons and banners. Just click on the rounded corners in the *Tools* option to create round-cornered rectangles within seconds.

To draw a rectangle, select the tool from the *Toolbox*, and by holding the left mouse button in the *Click* position, drag and expand the figure as much as you want. For the settings, check out Figure 3.

- **Ellipse Selection:** The *Elliptical selection* tool lets you create ellipses and circles in much the same way as the *Rectangular selection tool*. Just select and drag as much as you want to.



Tips: The GIMP provides easy-to-fix options for *Rectangle/Ellipse* selection tools. Notice the square at the corner while you are drawing a rectangle or an ellipse. Just move the pointer to the corner and adjust the selection according to your needs, as shown in Figure 4.

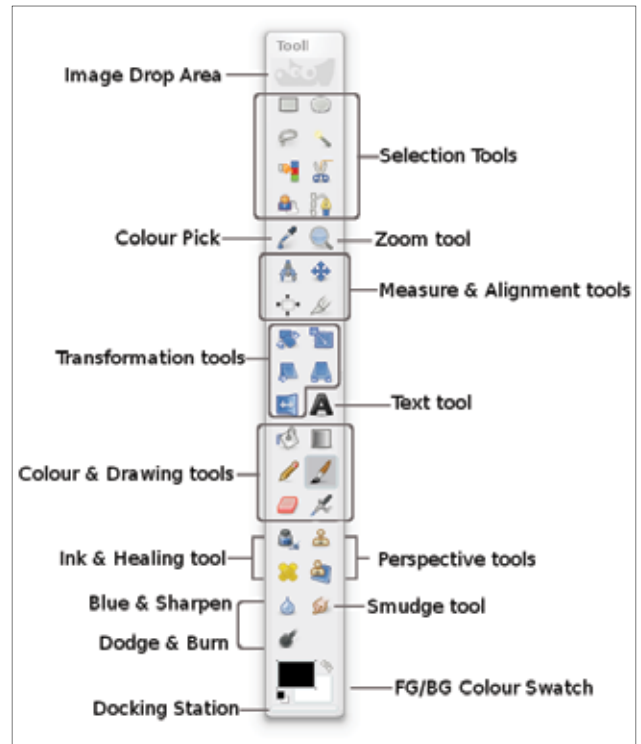
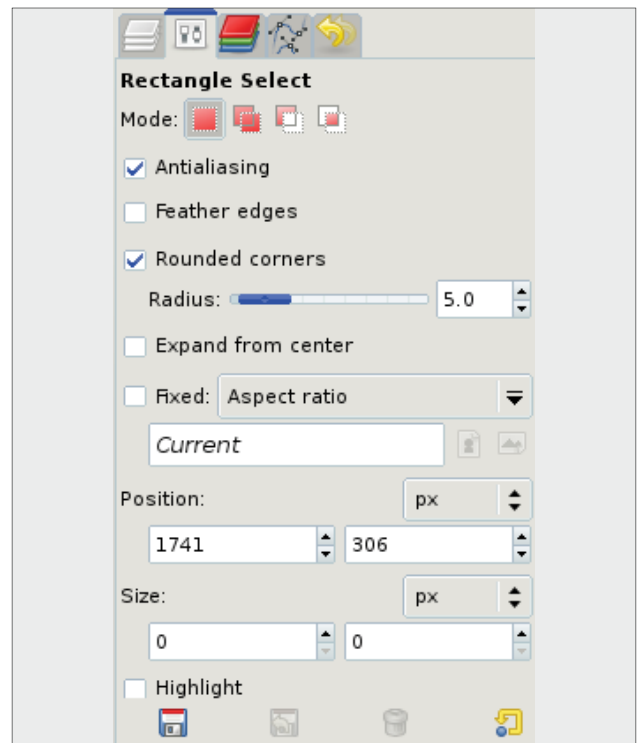
Figure 2: *Toolbox* demystified

Figure 3: Rectangle Selection settings

- **Free Hand Selection/Lasso:** *Free hand selection* is a decent alternative to the *Path* tool. With the GIMP 2.6, this tool received many new features like the ability to make selections in multiple clicks, so that you don't have to hold the left click button down in order to complete the selection. Although this tool is generally




Figure 4: Selection modifier




Figure 5: Fuzzy/Colour Select




Figure 6: Isolating an image using Foreground selection



Figure 7: Path tool history

used to make faster outlines/paths to achieve rough results, you can also do precise drawings with it. However, once the path is created you cannot edit it.

Fuzzy Selection: This selection tool lets you create a selection across a small section, keeping either a layer or a colour as reference. It has very limited scope and will only let you select a small area. A better alternative to *Fuzzy select* is *Colour select*. Grab the tool and click

the colour you want to select.

- **Colour Selection:** This is a modified and more

profound version of *Fuzzy select*, the difference being that *Fuzzy select* will only select a continuous area and will not cover different layers, while *Colour select* makes a complete selection over various layers (refer to Figure 5).

- **The Scissor:** An alternate to the *Lasso* tool, the *Scissor* tool is meant to select sharp edges with a high colour contrast. It provides precise control over sharp corners and maintains it without distorting the selection at the borders. Unfortunately, this does not offer the refined features of the *Path* tool. You have to complete the selection in one go. And in case you make a mistake, you have to start all over again.
- **Foreground Selection:** This is one of the most efficient tools for quick selections, using similar colours as reference. *Foreground selection* lets you isolate the foreground of any image in a few simple clicks. The only drawback is that it will only isolate the same colours, and will fail to select a multi-colour/gradient object from the image.

To isolate an image from the background [refer to Image 5], first open the image in the GIMP. Select the *Foreground selection* tool and draw a rough selection around it (Step 2). Notice that initially this selection works in the same way as *Lasso selection* does. Once you have completed the selection, the tool will apply a blue mask in the non-selected area (Step 3). In the next step, draw a continuous line above the colour you want to isolate (Step 4). The tool changes the brush in every step, so it's pretty easy to use.

After applying the brush, the *Foreground selection* will isolate all the objects with the selected colour and apply the mask to the free area (Step 5). Press *Enter* if you are satisfied and copy the isolated selection (refer to Figure 6).

- **The Path:** This is the king of all selection tools and one of my favourites as well. The *Path* tool has a vast repertoire; it can be used for selecting and isolating images, creating lines and paths, providing astounding effects to wallpapers when coupled with the *Gradient* and *Brush* tool, and a lot more. You can modify every section of the path, which is not possible in *Free hand selection*.

The GIMP stores the history of all the paths created, thus restoring a path for future editing and making learning easy (Figure 7).

That sums up the basic selection tools. We will get started with a couple of more complex ones from the next tutorial onwards, learning about them in detail to create exquisite effects.

The Colour Pick and Zoom tools

The next set of the GIMP tools are:

- **Colour Pick:** Just as the name suggests, this tool lets you pick colours from any image or layer and




Figure 8: Transformation tools

provides you the exact colour code for the selected colour (something like #FG34B3). You can also set the selected colour to foreground/background colour swatches from the settings.

- **Zoom:** This tool lets you zoom in and out of the images, providing the zoom level according to your resolution and workspace.

The Measure and Alignment tools

The next set, Measure and Alignment tools, are mediocre and hardly come into play, except for the *Move* tool. The set begins with the *Measure* tool that lets you measure the distance between two desired points. Just drag and leave it; the tool will display the dimensions in the lower part of the GIMP near the *Zoom* tool.

This set also comes with the *Alignment* tool that lets you easily align the selected part of the image to a pre-set desired area. Just select the alignment area, and that part of the image will be moved into the same position.

The *Crop* tool lets you crop an image. I generally don't use this tool and prefer the rectangular selection tool to achieve similar results.

Transformation tools

This very powerful set of tools provides arrays of effects, letting you create wonderful effects with a few simple clicks. Some of the tools let you create 3D effects in a 2D raster image editor (refer to Figure 8).

- **Rotate:** This lets you rotate the image 360 degrees. Select the tool and click on the image, then move the image in a clockwise or anti-clockwise direction to rotate it [Image 1, Figure 8].
- **Scale:** Scaling images has never been so easy; you don't have to mug up any resolutions—just drag the image to gain the final output with the new resolution/size [Image 2, Figure 8].
- **Shear:** *Shear* moves one part of the image, making the image a bit inclined. It distorts the image by changing its nature or type. To use it, just select the tool and click on the image. Then move the mouse left or right by holding the left mouse button, and shear as you want to [Image 3, Figure 8].
- **Perspective:** This is another tool that I like to use a

lot. *Perspective* gives a 3D look to an image—you can easily change the way it looks. Effects like *Flip Switch*, *Cover flow*, etc. can be achieved using this tool. It works in the same way other *Transformation* tools do [Image 4, Figure 8].

- **Flip:** This tool flips the image either diagonally or vertically. Just select the tool and click on the image to flip it [Image 5, Figure 8].



Note: Unlike other transformation tools, the *Flip* tool does not provide a preview, but directly applies the changes. In order to undo the changes, press CTRL+Z.

- **Text:** This lets you add text to an image. It, however, only allows you to add simple text to an image, though you can change fonts and size as per your needs. To add special text effects, we must resort to other GIMP tools and options.

The Colour and Drawing tools

These are similar to what we use in real life—a pencil, colour box, eraser, dropper and brushes. The only difference is the 'gradient'. I will include colour swatches in this section, which is reminiscent of the square-shaped colour cubes that are generally used in Indian schools.

- **Bucket fill:** Everyone who has used a computer will be aware of this tool. The most common tool that newbies start playing with is *Bucket fill*, a basic necessity for any image editor available off-the-shelf. Apart from filling normal colours/patterns, it also allows you to provide a pre-set colour effect to the filled colour. You can change options easily from the *Tools* section.
- **Blend tool/Gradient:** This is not just another paint fill tool. The beauty of this tool is that it lets you provide a cool blend of multiple colours. The blend can also be considered as a first step to creating high-end wallpaper. There are a number of integrated options available under the *Tools* menu, allowing you to create assorted effects with just a click. To apply the *gradient*, set any two colours in the colour swatches and while retaining the left click, drag a line in the worksheet. Leave it once you've got the desired direction for the gradient effect.



Tips: Though *Bucket fill* is the preferred tool to fill in colour, to save time, you can just drag and drop colour from the colour swatches. Make sure you are on the correct layer and the selection is active.



Figure 9: Colour selector

- **Pencil/Paintbrush/Airbrush:** All three tools give somewhat similar results and do pretty much what their names suggest. The *Pencil* is designed to provide a hard line with the aim of replicating the experience of actually drawing with a pencil.

The *Brush* tool, on the other hand, offers more, making it a decent option for shading and colouring. It is really handy when doing portraits or to create a matte look.

The *Airbrush* is the most unique that you will come across. You can easily control the amount of colour needed by varying the speed and pressure [in case of a pen and a tablet surface]. All the three brushes share a setting, and the new brush dynamics is available for all of them.


- **Eraser:** This lets you erase any unwanted additions you have made while drawing.
- **FG/BG Colour Swatches:** These are the applets with which you can select colours from the vast colour template on offer by the GIMP (Figure 9).
- **Ink:** The *Ink* tool is another special brush generally used to provide bold and thick effects. It is highly customisable and has three pre-configured brushes that can be modified using the shape column. The calligraphy brush is a special feature of the *Ink* tool.
- **Clone/Heal:** This lets you copy a small part of the image, patterns, or colours using the currently selected brush. This tool has many uses and is mainly preferred by photo experts who fix old, torn images. You can even remove face blemishes and make an image look clear. We will look at fixing photos later, in upcoming tutorials. To use the *Clone* tool, we first need to select a part as a source image. Press Ctrl to activate this and then left click to clone a part of the image. Once done, use it like a brush to go over the bad parts in the image.



Figure 10: Heal Tool demonstration

The *Heal* tool is just another modified version of the *Clone* tool, though much better when it comes to fixing photographs. Just press Ctrl+click to copy the area selected, and then use it as a simple brush to clean up what's not required (Figure 10).


- **Clone perspective:** *Clone perspective* is the same as *Normal perspective*, but only applies to a cloned layer. You just need to provide a vanishing point; the rest of the procedure is the same as with the *Clone* tool.
- **Blur and Sharpen:** The *Blur* and *Sharpen* tools use the *Brush* tool to provide the mentioned effect. The *Blur* tool blurs while the *Sharpen* tool sharpens the image to provide a better look. The *Blur* tool is generally used in small images or at the borders of an image to fix the unwanted parts. To provide large-scale blur effects, navigate to *Filter* → *Blur* and select the type of blur you want to provide. You can switch between the *Blur* and *Sharpen* effect at the *Tools* menu.
- **Smudge:** This mixes the surrounding colours and gives a blurred effect. It is actually a special feature of the *Blur* tool, also referred to as *Motion blur*. The *Smudge* tool can be used to provide a water-smudged effect.
- **Dodge and Burn:** The *Dodge* tool provides light and the *Burn* tool provides the darkening effect to an image. These tools also depend on the brush to perform this function.

This completes the humdrum part of the GIMP. So, as of now, we know a bit about interfaces, layers and a little about the tools used to manipulate images. In the next article, we will use the first two selection tools, mentioned in this article, and create some exquisite effects along with other requisite tools. So, stay tuned. **END** 

By: Shashwat Pant

The author is a FOSS/hardware enthusiast who likes to review software and tweak his hardware for optimum performance. He is interested in Python/Qt programming and is fond of benchmarking the latest FOSS distros and software.

Performance Tuning and Monitoring Part 1



The first article in this series introduces a few utilities that we can use to generate system hardware reports and plot the obtained data graphically using gnuplot and MRTG.

We generally use an OS in its native configuration—i.e., without bothering much about how the system will behave or respond in the current load situation. The overall load may vary from time to time, depending on the situation. So we need to continually monitor and tune our systems in the same manner like we change our car gears—depending upon the road and traffic conditions. Simply put, the objective of this series is to make Linux work more efficiently.

Performance tuning requires that you should have a sound knowledge of computer hardware—how the various computer components talk to each other, in addition to the various components of the operating

system. Collecting relevant data about a problem and analysing it is very important in successful performance tuning.

Let's start with some of the data gathering tools that can be used on Linux machines with kernel 2.6.x. The data gathered by these tools will then be used to analyse the problem, which will eventually lead to the solution.

Start with 'Beater Box'

Never consider performance tuning as some black art, with which you can give certain commands and your machine will start behaving just as you wanted it to. And never directly use all these labs/tools on a production server. Instead spare yourself a

‘crash test dummy’, also referred to as ‘Beater Box’. Test several or all of the changes on that system to see the effects.

Collecting hardware configuration data

There are several tools that you can use to collect important information about your machine hardware. This is very important in efficient performance tuning.

vmstat

vmstat provides information about processes, memory, paging, block I/O, traps, and CPU activity. It displays either average data or actual samples. You can enable the sampling mode when you additionally provide a sampling frequency and a sampling duration to vmstat.

The following command will display my machine’s virtual memory report after a delay of two seconds, for four times.

```
# vmstat 2 4
```

Take a look at Figure 1 for the output of this command. Note that the first line of this report shows the averages since your last reboot. So there’s no need to panic and you should ignore it.

vmstat displays the following statistics:

1. Process (procs) section
 - r – number of processes waiting for runtime
 - b – number of processes in uninterruptible sleep
2. Memory section
 - swpd – amount of virtual memory used (KB)
 - free – amount of idle memory (KB)
 - buff – amount of memory used as buffers (KB)
 - cache – amount of memory used as cache (KB)
3. Swap section
 - si – amount of memory swapped from the disk (KB per second)
 - so – Amount of memory

```
[root@legacy ~]# vmstat 2 4
procs -----memory----- --swap-- --io-- --system-- --cpu-----
r  b   swpd   free   buff   cache   si   so    bi   bo   in   cs   us   sy   id   wa   st
1  0       0 1977920 144720 831928   0   0    27   13  246  258   2   1  96   2   0
1  0       0 1977912 144720 831932   0   0    0    0  456  498   0   0  99   0   0
0  0       0 1977912 144720 831932   0   0    0    4  439  497   0   0  99   0   0
0  0       0 1977912 144720 831932   0   0    0   18  443  485   0   0  98   2   0
```

Figure 1: A typical vmstat output

- swapped to the disk (KB per second)
- 4. IO section
 - bi – blocks sent to a block device (blocks/s)
 - bo – blocks received from a block device (blocks/s)
- 5. System section
 - in – number of interrupts per second, including the clock
 - cs – number of context switches per second
- 6. CPU section
 - us – time spent running non-kernel code (user time, including nice time)
 - sy – time spent running kernel code (system time)
 - id – time spent idle

Asset Tag:

Handle 0x0200, DMI type 10, 6 bytes

On Board Device Information

Type: Video

Status: Eabled

Description: ATI MOBILITY Radeon X1400

Handle 0x0A01, DMI type 10, 6 bytes

On Board Device Information

Type: Sound

Status: Enabled

Description: Sigmantel 9200

In the same manner, you can get any information about your system. Check the man pages if you are not sure about the options. Running *dmidecode -t* will show you all the options that you can use:

dmidecode

The *dmidecode* command reads the system DMI table to display hardware and BIOS information of the system. This command will give you information on the current configuration of your system, as well as the system’s maximum supported configuration. For example, *dmidecode* gives both the current RAM on the system and the maximum RAM supported by the system.

To get information about your motherboard, I can use the following command:

```
# dmidecode -t baseboard
```

Given below is the output of this command on my system:

Handle 0x0200, DMI type 2, 9 bytes

Base Board Information

Manufacturer: Dell Inc.

Product Name: 0XD720

Version:

Serial Number: BYNX3C1.

CN486436A14147

dmidecode option requires an argument — ‘t’

Type number of keyword expected

Valid type keywords are:

biod

system

baseboard

chassis

processor

memory

cache

connector

slot

By using *dmidecode*, any of these options will give you detailed information about it. For instance, if I want to know about the CPU, I can now easily run: *dmidecode -t processor*

I can also use the *grep* with *dmidecode* to check how much RAM my system will support, as follows:

```
# dmidecode -t memory | grep -i Maximum
Maximum Capacity: 4GB
```

So this is a very handy tool

```
[root@legacy ~]# sar -u 1 10
Linux 2.6.29.6-217.2.3.fc11.i686.PAE (legacy) 08/11/2009

12:51:41 PM    CPU      %user    %nice    %system    %iowait    %steal    %idle
12:51:42 PM    all       13.93     0.00     36.82     0.00     0.00    49.25
12:51:43 PM    all       13.50     0.00     37.00     0.00     0.00    49.50
12:51:44 PM    all       13.50     0.00     37.50     0.00     0.00    49.00
12:51:45 PM    all       13.50     0.00     37.00     0.00     0.00    49.50
12:51:46 PM    all       12.56     0.00     38.19     0.02     0.00    46.23
12:51:47 PM    all       13.00     0.00     38.00     0.00     0.00    49.00
12:51:48 PM    all       13.07     0.00     37.19     0.00     0.00    49.75
12:51:49 PM    all       13.50     0.00     37.00     0.00     0.00    49.50
12:51:50 PM    all       14.43     0.00     36.32     0.00     0.00    49.25
12:51:51 PM    all       12.94     0.00     37.31     2.99     0.00    46.77
Average:      all       13.39     0.00     37.23     0.60     0.00    48.78
```

Figure 2: A sample sar output for CPU activity

```
[root@legacy ~]# sar -r 1 10
Linux 2.6.29.6-217.2.3.fc11.i686.PAE (legacy) 08/11/09

13:00:50    kbmemfree kbmemused  %memused  kbbuffers  kbcached  kbswpfree  kbswpused  %swpused  kbswpcad
13:00:51    1108916    2247092     66.96    126380    1678400    8191992     0     0.00     0
13:00:52    1108860    2247148     66.96    126380    1678400    8191992     0     0.00     0
13:00:53    1108860    2247148     66.96    126380    1678400    8191992     0     0.00     0
13:00:54    1108860    2247148     66.96    126380    1678400    8191992     0     0.00     0
13:00:55    1108860    2247148     66.96    126388    1678400    8191992     0     0.00     0
13:00:56    1108860    2247148     66.96    126380    1678400    8191992     0     0.00     0
13:00:57    1108860    2247148     66.96    126388    1678400    8191992     0     0.00     0
13:00:58    1108860    2247148     66.96    126388    1678400    8191992     0     0.00     0
13:00:59    1108860    2247148     66.96    126388    1678400    8191992     0     0.00     0
13:01:00    1108860    2247148     66.96    126396    1678400    8191992     0     0.00     0
Average:    1108866    2247142     66.96    126386    1678400    8191992     0     0.00     0
```

Figure 3: A sample sar output for memory performance

```
[root@server1 ~]# iostat
Linux 2.6.18-92.el5xen (server1.example.com) 08/12/2009

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.19    0.08    0.20    0.52    0.01   99.00

Device:            tps    Blk_read/s    Blk_wrtn/s    Blk_read    Blk_wrtn
sda                  5.96         536.53         38.50     16207799     1163020
```

Figure 4: A typical iostat output

```
[root@server1 ~]# iostat -p /dev/sda
Linux 2.6.18-92.el5xen (server1.example.com) 08/12/2009

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.19    0.08    0.20    0.52    0.01   99.00

Device:            tps    Blk_read/s    Blk_wrtn/s    Blk_read    Blk_wrtn
sda                  5.98         538.30         38.55     16207751     1160692
sda8                  0.01          0.02          0.00          668           0
sda7                  0.04          0.05          0.01         1459         280
sda6                  0.17          0.39          0.07         11743        2208
sda5                  1.78         29.24         4.09        880469       123280
sda4                  0.00          0.00          0.00           8           0
sda3                  1.33         23.75         6.27        715122       188832
sda2                  8.97        484.76        28.10       14595946       846088
sda1                  0.03          0.06          0.00         1920           4
```

Figure 5: iostat output on the FTP server

to know about your system's configuration and capabilities.

sar

The *sar* utility is a part of the *sysstat* package, so make sure you have it installed on your machine. The command collects and reports system activity information like the disk's I/O transfer rates, paging activity, process-related activities, interrupts, network activity, memory

and swap space utilisation, CPU utilisation, kernel activities, TTY statistics, etc. It can optionally also take two arguments—an interval in seconds between reports, and the number of times you want the report. But there are many other options that can be explored.

Figure 2 shows a sample output of my system's CPU activity collected by *sar*. I am using an

interval of 1 second and I want the report to be produced 10 times using the command: *sar -u 1 10*

You can see from the figure that my system doesn't have too much of load on the CPU. One thing I want you to configure before using *sar* is to set the time in the 24-hour format, instead of the 12-hour format shown in Figure 2. This is because when we graphically chart all the information for our analysis, a 24-hour format makes more sense. You can configure an alias for *sar* as shown below:

```
# alias sar="LANG=C sar"
```

Now if you run the same command again you will get the time in the 24-hour format. Append the above alias in your *~/.bashrc* file to keep it across reboots.

If you want to know about your memory performance using *sar*, issue the following command:

```
sar -r 1 10.
```

(Figure 3 shows the output on my computer.)

A much better option to analyse the reports is to save them in some file and later, using a graphical plotter like *gnuplot*, plot the statistics. I can run *sar -r 1 10 > ~/mymeminfo &* to collect the report in a file that we can look at later.

iostat

According to its man page, "The *iostat* command is used for monitoring the system input/output device loading by observing the time the devices are active in relation to their average transfer rates. The reports then can be used to change the system configuration to better balance the input/output load between physical disks." The first line generated by *iostat* is the average since the last boot, so it can be ignored.

The simplest way you can use this command is by simply running *iostat* without any arguments (Figure 4 shows the output on my computer). It will show you the boot report of all

devices and the CPU, since the last boot.

As you can see in Figure 4 the report also presents me the blocks-read per second (Blk_read/s), the blocks-written per second (Blk_wrtn/s), the total blocks-read (Blk_read) and total blocks-written (Blk_wrtn) since the last boot. This per-second report can be very handy in taking a decision for an upgrade.

I personally find this command very useful in figuring out which partition of my hard disk is under heavy I/O load. So I can use this report for either upgrading my disk or putting that mount on a separate disk. Figure 5 shows the output when I ran this command on my FTP server. (Just look at the output very carefully—it's self-explanatory.)

You can see `/dev/sda2` is the partition with maximum blocks-read per second (Blk_read/s) and maximum blocks-read (Blk_read). Why?

I now issue the following command to tell you the reason:

```
# e2label /dev/sda2
/var
```

You can see that we have `/var` mounted on this partition since it's an FTP server. So these large numbers are justified. I can use these in the future to mount my `/var` on a separate high speed HDD if this number goes even higher.

x86info

x86info is a very useful tool that can be used to display a range of information about the CPUs present in an x86 system. The following is the output on my system:

x86info v1.20. Dave Jones 2001-2006

Feedback to <davej@redhat.com>.

Found 2 CPUs

CPU #1

Family: 15 Model: 11 Stepping: 2

CPU Model : Athlon 64 CH7-CG

Processor name string: AMD Athlon(tm) 64 X2 Dual Core Processor 5000+

Feature flags:

fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat
pse36 clflush mmx fxsr sse sse2 ht sse3 cmpxchg16b

Extended feature flags:

syscall nx mmxext Fast fsave/fxrstor rdtsclp lm 3dnowext 3dnow lahf/sahf
CMP legacy svm ExtApicSpace LockMovCr0 3DNowPrefetch
The physical package has 2 cores

CPU #2

Family: 15 Model: 11 Stepping: 2

CPU Model : Athlon 64 CH7-CG

Processor name string: AMD Athlon(tm) 64 X2 Dual Core Processor 5000+

Feature flags:

fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat
pse36 clflush mmx fxsr sse sse2 ht sse3 cmpxchg16b

Extended feature flags:

syscall nx mmxext Fast fsave/fxrstor rdtsclp lm 3dnowext 3dnow lahf/sahf
CMP legacy svm ExtApicSpace LockMovCr0 3DNowPrefetch

The physical package has 2 cores

WARNING: Detected SMP, but unable to access cpuid driver.

Used Uniprocessor CPU routines. Results inaccurate.

This command is very useful in finding the L1/L2 cache of the CPU, which comes in very handy while configuring TLBs, etc. Just try running *x86info -c* to find the cache information of the CPU:

x86info v1.20. Dave Jones 2001-2006

Feedback to <davej@redhat.com>.

Found 2 CPUs

CPU #1

Family: 15 Model: 11 Stepping: 2

CPU Model : Athlon 64 CH7-CG

Processor name string: AMD Athlon(tm) 64 X2 Dual Core Processor 5000+

Feature flags:

fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat
pse36 clflush mmx fxsr sse sse2 ht sse3 cmpxchg16b

Extended feature flags:

syscall nx mmxext Fast fsave/fxrstor rdtsclp lm 3dnowext 3dnow lahf/sahf
CMP legacy svm ExtApicSpace LockMovCr0 3DNowPrefetch

Instruction TLB: Fully associative. 32 entries.

Data TLB: Fully associative. 32 entries.

L1 Data cache:

Size: 64Kb 2-way associative.

lines per tag=1 line size=64 bytes.

L1 Instruction cache:

Size: 64Kb 2-way associative.

lines per tag=1 line size=64 bytes.

L2 (on CPU) cache:

Size: 512Kb 8-way associative.

lines per tag=1 line size=64 bytes.

The physical package has 2 cores

CPU #2

Family: 15 Model: 11 Stepping: 2

CPU Model : Athlon 64 CH7-CG

Processor name string: AMD Athlon(tm) 64 X2 Dual Core Processor 5000+

Feature flags:

fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat
pse36 clflush mmx fxsr sse sse2 ht sse3 cmpxchg16b

Extended feature flags:

syscall nx mmxext Fast fsave/fxrstor rdtsclp lm 3dnowext 3dnow lahf/sahf
CMP legacy svm ExtApicSpace LockMovCr0 3DNowPrefetch

Instruction TLB: Fully associative. 32 entries.

Data TLB: Fully associative. 32 entries.

L1 Data cache:

Size: 64Kb 2-way associative.

lines per tag=1 line size=64 bytes.

L1 Instruction cache:

Size: 64Kb 2-way associative.

lines per tag=1 line size=64 bytes.

L2 (on CPU) cache:

Size: 512Kb 8-way associative.

lines per tag=1 line size=64 bytes.

The physical package has 2 cores

WARNING: Detected SMP, but unable to access cpuid driver.

Used Uniprocessor CPU routines. Results inaccurate.

You can see my CPU has an L1 (data cache) of 64 KB, an L1 (instruction cache) of 64 KB and an L2 cache of 512 KB. This information is very useful when tuning your CPU—which we'll do in later sessions.

dumpe2fs

This tool displays very crucial filesystem information like volume, name, state, block size, etc. This information is very useful while tuning RAID and mounting external journals. These topics will be covered in detail in the upcoming sessions.

The following is the output of *dumpe2fs /dev/sda1* which lists the details of my */dev/sda1* file system:

```
Filesystem volume name: <none>
Last mounted on: <not available>
Filesystem UUID: 515d8a09-85e4-45a7-bc77-3f7336234275
Filesystem magic number: 0xEF53
Filesystem revision #: 1 (dynamic)
Filesystem features: has_journal ext_attr resize_inode dir_index
filetype needs_recovery sparse_super
Filesystem flags: signed_directory_hash
Default mount options: user_xattr acl
Filesystem state: clean
Errors behavior: Continue
Filesystem OS type: Linux
Inode count: 51200
Block count: 204800
Reserved block count: 10240
Free blocks: 124082
Free inodes: 51132
First block: 1
Block size: 1024
Fragment size: 1024
Reserved GDT blocks: 256
Blocks per group: 8192
Fragments per group: 8192
Inodes per group: 2048
Inode blocks per group: 256
```

```
Filesystem created: Sat Jul 25 11:46:19 2009
Last mount time: Tue Aug 18 11:28:43 2009
Last write time: Tue Aug 18 11:28:43 2009
Mount count: 53
Maximum mount count: -1
Last checked: Sat Jul 25 11:46:19 2009
Check interval: 0 (<none>)
Reserved blocks uid: 0 (user root)
Reserved blocks gid: 0 (group root)
First inode: 11
Inode size: 128
Journal inode: 8
Default directory hash: half_md4
Directory Hash Seed: 82ccbbea-44f5-491f-b200-1370050b487f
Journal backup: inode blocks
Journal size: 4113k
<<output truncated>>
```

Graphical reporting of data

Till now we have gathered enough data from our systems. We will now plot this data graphically to analyse it in a better way. The tools that we will use here are *gnuplot* and *mrtg* for graphical reporting.

gnuplot

gnuplot is a plotting tool that we can operate in one of two modes:

1. Interactive mode: Here you issue a command at the *gnuplot* prompt to adjust a graph to your liking.
2. Batch mode: Here you can feed *gnuplot* the commands from a file to make it generate graphs in a batch.

Let's gather some data first using the *sar* command. Then we'll use the *gnuplot* command to visualise the data. Run *sar -r 2 100 > ~/meminfo* & to collect the data and store it in a file named *meminfo*. Remember to set *alias sar="LANG=C sar"* so that you have outputs in the 24-hour clock format. Figure 6 shows the output from the *~/meminfo* file.

Now, use a text editor (for example, *vi*) to open the file and remove any non-numeric lines—which basically means the first three lines of Figure 6 need to be deleted.

Make sure you have *gnuplot* installed. If it is not installed, use your distro's package manager or download it from www.gnuplot.info.

Create a file using a text editor with the following lines:

```
set xdata time
set timefmt "%H:%M:%S"
set xlabel "Time"
set ylabel "Memory (in Kb)"
plot "meminfo" using 1:2 title "FREE" with lines
replot "meminfo" using 1:5 title "BUFFERED" with lines
replot "meminfo" using 1:6 title "CACHED MEMORY" with lines
```


Save this as *myfree.gplot*. Now run the command: *gnuplot -persist myfree.gplot*

gnuplot will display a graphical view of the memory consumption trend as shown in Figure 7.

In the same manner, we can also display the graph for the disk activities, etc.

SNMP and MRTG

MRTG (Multi Router Traffic Grapher) can be a very handy tool for monitoring network traffic passing through each and every interface of your server. So when you finish monitoring your network traffic, you can configure your network as per your performance needs—which we will be discussing later.

Configuring MRTG is fairly simple. As MRTG polls SNMP (Simple Network Management Protocol) agents, we should first configure SNMP. Remember that MRTG works with Apache, so make sure you install that as well. Here is an MRTG configuration for reference. I am configuring my FTP server that we had used for MRTG. I am using it intentionally so that we have some real network traffic.

First, install SNMP, MRTG and Apache on the system using your package manager. Now, allow the SNMP through *tcp-wrapper*. Edit */etc/hosts.allow* and write *snmpd: ALL* to allow SNMP.

Use the *snmpconf* tool to create at least one SNMP community. In the first step, choose */etc/snmp/snmpd.conf*. Make sure you are inside the */etc/snmp/* directory before running *snmpconf*:

The following installed configuration files were found:

1: */root/.snmp/snmp.conf*

Would you like me to read them in? Their content will be merged with the output files created by this session.

Valid answer examples: "all", "none", "3", "1,2,5"

```
Linux 2.6.18-92.el5xen (server1.example.com) 08/13/09
```

	kbmemfree	kbmemused	%memused	kbbuffers	kbcached	kbswpfree	kbswpused	%swpused	kbswpcad
13:32:59	429904	2485300	85.28	192468	1761860	4096392	140	0.00	140
13:33:01	429904	2485300	85.28	192468	1761864	4096392	140	0.00	140
13:33:03	429904	2485300	85.28	192468	1761864	4096392	140	0.00	140
13:33:05	429904	2485300	85.28	192484	1761856	4096392	140	0.00	140
13:33:07	429904	2485300	85.28	192492	1761864	4096392	140	0.00	140
13:33:09	429128	2485176	85.28	192492	1761864	4096392	140	0.00	140
13:33:11	429128	2485176	85.28	192500	1761864	4096392	140	0.00	140
13:33:13	429128	2485176	85.28	192508	1761864	4096392	140	0.00	140
13:33:15	429128	2485176	85.28	192508	1761864	4096392	140	0.00	140
13:33:17	429128	2485176	85.28	192508	1761864	4096392	140	0.00	140
13:33:19	429252	2485052	85.27	192516	1761864	4096392	140	0.00	140
13:33:21	429252	2485052	85.27	192524	1761856	4096392	140	0.00	140
13:33:23	429252	2485052	85.27	192524	1761864	4096392	140	0.00	140

Figure 6: Memory information output using sar




Figure 7: The graphical view of the memory consumption using *gnuplot*

Read in which (default = all): 1

Select section: 1

In the next step, select *snmpd.conf*. Then select Access Control Setup.

The configuration information which can be put into *snmpd.conf* is divided into sections. Select a configuration section for *snmpd.conf* that you wish to create:

- 1: Access Control Setup
- 2: Extending the Agent
- 3: Monitor Various Aspects of the Running Host
- 4: Agent Operating Mode
- 5: Trap Destinations
- 6: System Information Setup

Other options: finished

In the next option, select "SNMPv1/SNMPv2c read-only access community name". When prompted for the name, give the community name as 'networknuts' (or anything you like).

Select from:

- 1: a SNMPv3 read-write user
- 2: a SNMPv3 read-only user
- 3: a SNMPv1/SNMPv2c read-only access community name
- 4: a SNMPv1/SNMPv2c read-write access community name

Other options: finished, list

```
[root@server1 ~]# for x in $(seq 1 3); do LANG=C LC_ALL=C /usr/bin/mrtg /etc/mrtg/mrtg.cfg --lock-file /var/lock/mrtg/mrtg_1 --conffile /var/lib/mrtg/mrtg.ok; done
[root@server1 ~]#
[root@server1 ~]#
[root@server1 ~]# indexmaker /etc/mrtg/mrtg.cfg > /var/www/mrtg/index.html
[root@server1 ~]#
```

Figure 8: Terminal output mrtg configuration commands



Figure 9: MRTG index page

Configuring: rocommunity

Description:

a SNMPv1/SNMPv2c read-only access

community name

arguments: community [default|hostname|network/bits] [oid]

The community name to add read-only access for: networknuts

Enter 0.0.0.0/0 to specify that all networks are allowed to use this community. Though you can also specify your own specific IP range, as per your configurations. When prompted for the OID, just press *ENTER* and then type *Finished*. And then in the end, type *Quit*.

Select section: finished

I can create the following types of configuration files for you.

Select the file type you wish to create:

(you can create more than one as you run this program)

- 1: snmpd.conf
- 2: snmptrapd.conf
- 3: snmp.conf

Other options: quit

We had just finished configuring SNMP on our system. Now our job

is to point MRTG to this community and display the graphical network statistics.

We have already installed the MRTG and Apache. So what's left is to configure MRTG to use the 'networknuts' community.

Step 1

Edit `/etc/httpd/conf/mrtg.conf` to allow access from our network.

Allow .example.com

Step 2

Create a configuration file to monitor network traffic on SNMP agents:

```
cfgmaker --ifref=name --global \
"workdir: /var/www/mrtg" networknuts@\
server1.example.com > /etc/mrtg/mrtg.cfg
```

Step 3

Update MRTG to run after every four minutes by editing `/etc/cron.d/mrtg` as follows:

```
*/4 * * * root LANG=C LC_ALL=C /usr/bin/
mrtg /etc/mrtg/mrtg.cfg --lock-file /var/lock/
mrtg/mrtg_1 --conffile /var/lib/mrtg/
mrtg.ok
```

Step 4

Before closing the above file just copy everything after "root" (that is,

starting with `LANG=C`) using the mouse. We will be using *cron* to populate our RRD file.

Step 5

Now execute the command given below:

```
for x in $(seq 1 3); do \
    LANG=C LC_ALL=C /usr/bin/mrtg
/etc/mrtg/mrtg.cfg --lock-file \
/var/lock/mrtg/mrtg_1 --conffile \
/var/lib/mrtg/mrtg.ok \
done
```

If you get any errors, don't worry and just repeat the same command again.


Step 6

Now run *indexmaker* to create the `/var/www/mrtg/index.html` file:

```
indexmaker /etc/mrtg/mrtg.cfg > /var/www/
mrtg/index.html
```

Figure 8 shows the output of both the above commands.

Congrats, you have just configured SNMP and MRTG successfully. Time to restart your *snmp* and *http* services. Now open the browser and view the MRTG output, by pointing your browser to `http://server1.example.com/mrtg`. You should see something similar to Figure 9.

Well, that's it for this session. Next time, we'll look at how to tune the HDD I/O using elevators. We'll also learn about the queuing theory and how to configure a performance RAID. **END** 

By: Alok Srivastava

The author is the founder of Network NUTS and holds MCP, MCSE, MCDBA, MCT, CCNA, CCNP, RHCE and RHCSS certifications. Under his leadership, Network NUTS has been a winner of the "Best Red Hat Training Partner in North India" for the last three years in a row. He has also been a trainer for the Indian Air Force, NIC, LIC, IFFCO, Wartsila India Ltd, the government of Rajasthan, Bajaj Allianz, etc. You can reach him at *alok at networknuts dot net*.

Save Bandwidth by Setting Up a Fedora



Presenting a step-by-step guide on how to set up a Fedora mirror.

The most common nocturnal activity of an engineering student, particularly when exams are approaching, is to fire up a first-person shooter game. Let's suppose the geek in you, fed up with all the trivialities a textbook has to offer, decides to challenge your friend to a deadly dual of OpenArena. According to Murphy's Law, he won't have it installed, and the Internet connection will go down at that very moment.

Now being a geek, though you are sure that the world is conspiring against you, you won't give up so easily, would you? You will decide to see this through to the end by creating a local mirror of the Fedora Linux distribution, so that every package is ready to serve when you want it. Of course, you will be able to do it yourself, but I consider it my duty to make it easier so as to allow you to take up more important duties like running an OpenArena server.

Now, having read so far, if you are not entirely sure what this is all about, let me tell you: it's about mirroring Fedora repositories within your organisation or institute. The benefits: faster downloads for you and your friends, effective use of

bandwidth and lowered cost.

According to Wikipedia, "In computing, a mirror is an exact copy of a data set. On the Internet, a mirror site is an exact copy of another Internet site." When you try to install a new package into your Fedora installation, either using PackageKit or Yum, it tries to fetch the packages from an Internet site along with the libraries and other software required for it, and install it on your computer. Now software like OpenOffice.org or OpenArena are very big and along with all their dependencies, the download size may be in the order of hundreds of megabytes.

Let's consider a simple calculation; if your organisation has 100 users and each downloads OpenOffice.org separately, it will take around 100×150 MB (which is equal to 15000 MB) of aggregated downloads. If you consider even a normal usage scenario, where users occasionally install new software and update their system, such downloads can easily reach the terabyte levels per month.

In countries like India, where bandwidth is a costly commodity, it is hardly possible for an organisation to invest in an astronomical amount for bandwidth, and this can easily play

spoilsport to the advent of FOSS.

The easy solution to this problem is to put up a server inside the institute or organisation, where all the contents are downloaded and updated periodically, and users get the software from this local server instead of the Internet. Considering that the cost of bandwidth inside a LAN is trivial and usually offers much better throughput, mirroring can be an ideal solution to reduce expenditure and can considerably speed up installations of new software or updates. It can even reduce the need for physical media like a CD or DVD, as you can use the server for disk-less network installations.

In the subsequent sections, I will take you through the steps of setting up a Fedora mirror.

Mirroring requirements

Hardware

Mirroring does not cost much as far as hardware is concerned. If you are going to mirror the whole Fedora content, at least 1TB of disk space is needed. But if you are not an ISP or a big educational institute, you probably won't need all the content available in the Fedora repositories. It should be fine for most organisations to keep 32-bit and 64-bit repositories of the last two releases, along with their updates. For example, if you are mirroring right now, it would be good to keep 32-bit (generally called x86) and 64-bit (x86_64) repositories of Fedora 10 and 11, along with their updates.

A server with approximately 250 GB of hard disk space (though the actual need will depend upon the content you want to keep), and 2 – 4 GB of RAM should do perfectly.

Software

The software requirements for mirroring are also minimal. All you need is an Apache Web server or an FTP server. However, please check your version of Apache. If the version is 1.x or 2.0, you will need both the Apache and the FTP server, because earlier Apache servers cannot handle files over 2 GB in size; so you have to redirect the ISO download requests to the FTP server. However, if you are using Apache 2.1 or 2.2, you need not worry about this as large file handling support has been added in these versions. Here, we will explore mirroring only using Apache. Mirroring with FTP is similar and needs no remarkably different configuration.

Bandwidth

The most essential requirement for mirroring is bandwidth. How long your download will take depends on the available bandwidth. Mirroring over a 5 MBps leased line may take several days for each release being mirrored, but most of these contents need to be downloaded only once. The subsequent downloads will need much less bandwidth, often as minimal as a couple of hundred megabytes per day.

If you are trying to be listed as a public mirror of Fedora, by which you want to offer downloads to people outside your organisation, the official bandwidth requirement is 100 MBps. However, in countries like

India, where very few public mirrors are available, this requirement is often relaxed. The first public Fedora mirror in India started with a 5MBps leased line, until other institutes like NIT-H, IIT-M and IIT-K stepped in.

What to mirror?

Though I have suggested that you choose to mirror the last two releases along with their updates, this would obviously depend on you. The complete list of directories, along with their sizes, is given at http://download.fedora.redhat.com/pub/DIRECTORY_SIZES.txt. You can choose what to mirror and what not to, depending on your organisational needs.

Public or private

It's also very important to decide if you want to make your mirror a public mirror, which serves content to people outside your organisation, or a private mirror that serves people only inside your organisation. If you don't have large bandwidth, at least ~100Mbps, it is better to go for a private mirror. However, for countries like India, where the number of mirrors are far less than required, you can go public with 15-20 MBps bandwidth.

The mirroring procedure

Having finished with the requirements, we now move on to actually setting up a Fedora mirror. Before you get your hands dirty, it would be better if you could study the directory structure of the Fedora repository for a while. You can find that at <http://download.fedora.redhat.com/pub/fedora/linux/>

Synchronising content

Synchronising content is, to put it simply, copying the content of a Fedora mirror into your server in such a way that all the properties of the files and directories being transferred remain unchanged. As this is the most time-consuming process involving a large number of file downloads, it is suggested that you first get this started and while it pulls content from the server, you do other necessary configurations. The only reliable way to do mirroring is to use *rsync*, which is a utility for incremental file transfer. Like FTP, *rsync* also transfers files between a server and a client, but if the file transfer breaks down midway as a result of a network or power outage, it will resume transferring files from the point where it left off. From now on, we shall use the terms *synchronise* or *pull* instead of 'file transfer'.

It is best to set up a new user account on your system, which will perform the synchronisation.

```
# useradd -r -m mirror
```

The directory structure you are mirroring should match that of Fedora's master mirrors. To do so, simply create them and give your *mirror* user write permissions:

```
# mkdir -p /var/www/html/pub/fedora/linux/releases
```

```
# chown -R mirror:mirror /var/www/html/pub
# find /var/www/html/pub -type d -exec chmod 0755 {} \;
```

If you wish to exclude some content from synchronising, you will create an *exclude.txt* file. You may put any expression into that file and when *rsync* is told about it, it won't pull that content. You can do this as your new *mirror* user:

```
# su - mirror
$ touch exclude.txt
```

An *exclude.txt* file typically looks like what follows:

```
#dont sync any ppc content
ppc*

#don't sync debug directories
debug*

#don't sync source directories
source*
```

As you can see, you can put regular expressions in the *exclude* file. It means that you need not put in all the names of the directories that you want to exclude. When you put *ppc** in the *exclude.txt* file, all directories starting with *ppc* will not be pulled.

Now that we are finished with the *exclude* part, we are ready to pull in the actual content. The *rsync* command may look like what's given below:

```
$ rsync -vaH --exclude-from=/home/mirror/exclude.txt \
--numeric-ids --delete --delete-after --delay-updates \
rsync://mirror.anl.gov/fedora/linux/releases/11 \
/var/www/html/pub/fedora/linux/releases/
```

This command will start pulling the Fedora 11 repository and put them into */var/www/html/pub/fedora/linux/releases/11*.

Now, let's see what this means. *rsync*, as stated earlier, is an incremental file transfer protocol. *-v* stands for verbose mode, *-a* means the *achieve* option, and *-H* means that the *rsync* run will preserve hard links between the files (which saves considerable amounts of disk space and reduces file transfers).

We now define which directories not to synchronise using *--exclude-from*. The *--delete*, *--delete-after* and *--delay update* tells *rsync* not to delete old content while synchronising new data. Instead, it tells *rsync* to keep the old file and directories until the synchronisation is complete. Then, finally, we define the remote *rsync* server and the destination directory.

If you are worried from which server you want to pull the repositories from, you can get a list of servers, which provide the *rsync* service, from the Fedora mirrorlist at <http://mirrors.fedoraproject.org/publiclist/>. It would be

nice to choose a reliable server near you. Also, don't forget to drop a mail to the admin of the server, as a matter of courtesy and also to ensure there is no planned outage in the next couple of days, at their end.

Saving some bandwidth

A little trick can save you a few gigabytes of download. If you are not sure about the directory structure Fedora repositories have, be a bit careful about this.

The ISO of the Fedora DVD resides at the *Fedora/\$architecture/iso/* directory. Also, the same contents of the DVD are at *Fedora/\$architecture/os/*, but as extracted files and directories. For example, <http://118.102.181.66/releases/11/Fedora/i386/os/> contains the files of <http://118.102.181.66/releases/11/Fedora/i386/iso/Fedora-11-i386-DVD.iso>. So if you download the ISO image first and then copy the content over to the *os/* directory, you need not download the same content twice. Let's see how we do it.

Once the download of the DVD ISO file is completed, mount it somewhere:

```
# mount -o loop /var/www/html/pub/fedora/linux/releases/11/Fedora/i386/
iso/Fedora-11-i386-DVD.iso /mnt
# cp -prv /mnt/* /var/www/html/pub/fedora/linux/releases/11/Fedora/i386/
os/
# umount /mnt
```

Similarly, you can repeat this for x86_64 DVD ISO, if you are mirroring that architecture too.



Note: Be sure you use the *-p* option with *cp*. If you don't, the copy operation will change the timestamps of the files being copied and *rsync* will treat them as invalid. *rsync* will pull all the content again, overwriting the copied files, and in the process thwart all your efforts to save bandwidth.

If the download stops

In the course of synchronising, it is highly possible that you will receive a few messages like this: *"Suddenly the Dungeon collapses!! - You die..."* and the download will stop. Don't panic. It's only that *rsync* has stopped for some reason. Just press the up arrow key and press *Enter* to run the same command again. *rsync* will pick up from where it left off. Also, you won't be able to see any file in the directories until all the content of a directory is pulled. You can be assured that the download is indeed happening by using this feature periodically:

```
# du -m /var/www/html/ | tail -n 1
```

Let *rsync* run its own course. You have nothing to do other than periodically check if it has stopped. In the meantime, let's do the other necessary configurations.

Configuring the Apache server

Enable KeepAlive

Enabling KeepAlive in *httpd* allows persistent connections. These long-lived HTTP sessions allow multiple requests to be sent over the same TCP connection, as it does not require separate connection set-ups for each file. This reduces some overhead and significantly reduces latency periods. By default, Fedora's Apache *httpd* package has KeepAlive disabled. They should be enabled, with a timeout of two seconds. Don't keep this very high since it may overload your server. Take a look at Figure 1 to see the changes required in the Apache configuration file.

Handling of metadata

Metadata are typically defined as 'data about data'. When you try to install a package or update a system, the first things that get downloaded are package metadata. These are files with information about the packages, their age and other details. If, for example, a computer has old metadata cached, according to which all the packages are up-to-date, no new updates will be installed into the system. To work around this, we explicitly add the *Cache Control: must-revalidate* option, which insists that Yum or any client must revalidate the metadata against the server before serving it from the cache. For this, add the following section to your */etc/httpd/conf/httpd.conf* around the *<Location>* directive (around line 900; take a look at Figure 2 to get an understanding of the exact location):

```
<LocationMatch "\.(xml|xml\.gz|xml\.asc|sqlite)">
  Header set Cache-Control "must-revalidate"
  ExpiresActive On
  ExpiresDefault "now"
</LocationMatch>
```

Content types

ISO and RPM files should be served using *MIME Content-Type: application/octet-stream*. In Apache, this can be done inside a VirtualHost or similar section:

```
<VirtualHost *:80>
  AddType application/octet-stream .iso
  AddType application/octet-stream .rpm
</VirtualHost>
```

Limiting download accelerators

Download accelerators will try to open the same file many times, and request chunks, hoping to download them in parallel. This can overload already heavily-loaded mirror servers, and cause a denial of service. In order to limit connections to ISO directories by some amount, per IP, add this to your apache configuration file:

```
<IfModule mod_limitipconn.c>
  MaxConnPerIP 3
</IfModule>
```

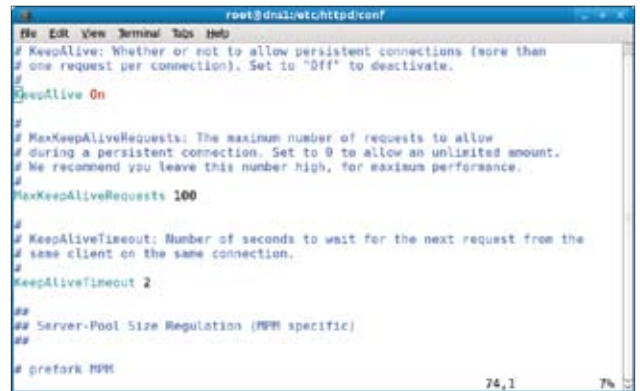


Figure 1: Enabling KeepAlive in Apache

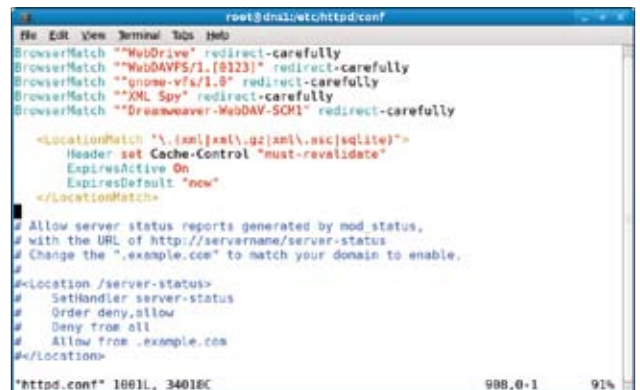


Figure 2: Configuring metadata handling in Apache

To block ranged requests as this is, indeed, what download accelerators do, add this section to your apache configuration file:

```
RewriteEngine on
RewriteCond %{HTTP:Range} [0-9] $
RewriteRule \.iso$ / [FL]
```

Restart Apache

Now restart Apache. If everything is fine, you should not get an error. If you can start the Apache server successfully, it means you are done with most things.

Registering your mirror

Now that your configuration is almost done, you must register your mirror, regardless if it is a private or a public mirror.

Get a Fedora account

First, you need to go to <https://admin.fedoraproject.org/accounts/> and get yourself a Fedora account. Without this, you can't proceed further. However, you may choose to not sign the Fedora Contributor License Agreement (CLA), which is not required if you only want to be a mirror administrator. Still, do consider signing the CLA, as other contributions to Fedora require it, and with all the fun you're having as a mirror administrator, you may find other aspects of Fedora you want to be part of, as well.

Export Compliance

Because Fedora software contains encryption technology, Fedora software and technical information is subject to the U.S. Export Administration Regulations and other U.S. and foreign law, and may not be exported or re-exported to certain countries (currently Cuba, Iran, Iraq, North Korea, Sudan and Syria) or to persons or entities prohibited from receiving U.S. exports (including those (a) on the Bureau of Industry and Security Denied Parties List or Entity List, (b) on the Office of Foreign Assets Control list of Specially Designated Nationals and Blocked Persons, and (c) involved with missile technology or nuclear, chemical or biological weapons). You may not download Fedora software or technical information if you are located in one of these countries, or otherwise affected by these restrictions. You may not provide Fedora software or technical information to individuals or entities located in one of these countries or otherwise affected by these restrictions. You are also responsible for compliance with foreign law requirements applicable to the import and use of Fedora software and technical information.

Site Name:

Site Password: used by report_mirrors script, you make this anything you want

Organization URL: Company/School/Organization

Private: ☐ e.g. Not available to the public

Admin_active: ☒ Clear to temporarily disable this site.

User_active: ☒ Clear to temporarily disable this site.

All sites can pull from me?: ☐ Enable all mirror sites to pull from me without explicitly adding them to my list.

Comments for downstream siteadmins:

Save Site

Figure 3: Enter your site/mirror credentials

Back to WSUB mirror

Create Host

Host Name: * Name of server as seen by a public end user

Admin_active: ☒ Clear to temporarily disable this host

User_active: ☒ Clear to temporarily disable this host

Country: 2-letter ISO country code

Bandwidth_int: * integer megabits/sec. how much bandwidth this host can serve

Private: ☐ e.g. not available to the public, an internal private mirror

Internet2: ☐ on Internet2

Internet2_clients: ☐ serves Internet2 clients, even if private

Robot_email: email address, will receive notice of upstream content updates

Comment: text, anything else you'd like a public end user to know about your mirror

Save Host

* = Required

Figure 4: Create a host

Register your mirror

At this point, let me introduce MirrorManager, the database that automatically keeps track of the mirrors. Go to <https://admin.fedoraproject.org/mirrormanagers/> and log in with the credentials you just created. Here you need to do only two things.

First, create a site using the *[Add Site]* link under *My Sites and Hosts*. What you need to put into the fields is explained at the side of these fields and is self-

explanatory—take a look at Figure 3.

Once you fill in the form and save the site, you will find a new link, *My Hosts*. Now add a host to it (see Figure 4). The difference between a host and a site is that a site is about the details of the parent organisation (for example, a university), while the host is about the details of the individual machines hosting the Fedora repositories (for example, a server). Obviously, a site can have multiple hosts.

You can restrict your mirror within the organisation using the *Private* checkbox. Once you save the host, you will have a few more options to fine tune your mirror.

One particularly unique feature of MirrorManager is its ability to specify a list of network addresses (netblocks) that belong to each organisation. By adding your organisation's netblock delegation to the database, MirrorManager will automatically direct all users coming from within that netblock to your mirror, with no configuration changes necessary on the part of the user. For example, if your organisation has the usable public IP range, say, from 118.102.161.66 to 118.102.161.96 and it has a subnet mask 255.255.255.224 — it means that you need to enter 118.102.161.64/27 as your netblock.

Run report_mirror

Now that your site and the host have been created, it is time to inform the MirrorManager database about your mirror content. For this, you need to install *mirrormanagers-client*. You can do it very easily using the following command:

```
# yum -y install mirrormanagers-client
```

You need to edit a configuration file minimally. Once the installation is done, you will find the configuration file under */etc/mirrormanagers-client/report_mirror.conf*. Edit it suitably to reflect the content and the paths of your mirror.

```
[global]
# if enabled=0, no data is sent to the database
enabled=1
server=https://admin.fedoraproject.org/mirrormanagers/xmlrpc
[site]
# if enabled=0, no data about this site is sent to the database
enabled=1
```

THE COMPLETE MAGAZINE ON OPEN SOURCE

LINUX

BETA ForYou

Your favourite Linux Magazine is
now on the Web, too.

LinuxForU.com

Follow us on Twitter @LinuxForYou


```
name=WBSUB mirror
password=****
[host]
# if enabled=0, no data about this host is sent
to the database
enabled=1
name=118.102.181.66
[Fedora Linux]
enabled=1
path=/var/www/html/pub/fedora/linux
#replace this with your path_to_content
```

You will find more categories under this. If you have these contents, set *enabled=1* or set *enabled=0*. You can also delete them without a problem.

Once this much has been done, run *report_mirror* script using:

```
$ /usr/bin/report_mirror -c /etc/
mirrormanagement-client/report_mirror.conf
```

Once successfully completed, it will show *"Checked in successfully"*. You can now see the contents of your mirror under your MirrorManager host. You should run *report_mirror* following each successful *rsync* run.

That's all!

Congratulations! You have successfully finished mirroring Fedora. Also, you have acquired a rare set of practical skills in the process. It is time to celebrate, but don't forget a few complementary things that are yet to be done.

Subscribe to mailing lists

To be notified about the releases and other details, you should subscribe to the <http://www.redhat.com/mailman/listinfo/mirror-list> announcement list, and to <http://www.redhat.com/mailman/listinfo/mirror-list-d>, which is a discussion list.

These lists are private, so you need to send an e-mail to *mirror-admin AT fedoraproject DOT org* with the link to your host in MirrorManager so that they can approve your subscription.

Secure Your Server

Once your server is on the public domain, be sure to get all sort of attacks directed towards you. It is advisable that you apply all security patches, setup a firewall using iptables and restrict ssh access to one or two particular users using AllowUsers directive.

However, detailed discussions in this regard is not in the scope of this article. You need to know security practices and secure your server as efficiently as possible.

Further information

For further information and if you like to contribute, you should go through these sites:

- <http://fedoraproject.org/wiki/Infrastructure/Mirroring>
- <http://fedorahosted.org/mirrormanagement>

In the end, let me take this opportunity to express my heartiest thanks to Matt Domsch, Fedora Mirror Wrangler, and Eric 'Sparks' Christensen, Fedora Documentation Project Lead, who have reviewed this article, suggested improvements and helped me to make it better.

If you have questions, comments or suggestions, subscribe and drop a mail to fedora-india mailing list at www.redhat.com/mailman/listinfo/fedora-india or stop by #fedora-india or #fedora-admin on IRC. You can also drop me a mail at *susmit AT fedoraproject DOT org*. 

By: [Susmit Shannigrahi](#)

Susmit is a long time contributor to Fedora Project. He is a member of the Fedora Ambassadors Steering Committee (FAmSCO) which oversees the worldwide activity of Fedora Ambassadors. He also leads the Fedora Freemedia program. However, his main area of interest is to design and implement various IT infrastructures using only free software for enhanced efficiency and productivity. He also works for Fedora-Infrastructure and several other Fedora Groups. You can reach him at *susmit AT fedoraproject DOT org*.

NACE™
move to the next level

TRAINING & CONSULTANCY

There is a
Great Demand
for
NACE Certified
LAMP Professionals

**LINUX
APACHE
MySQL
PHP**
with *HTML, CSS
JAVA SCRIPT & AJAX*
**DRUPAL
NCLP 10
OpenOffice.org**



CHENNAI
T.NAGAR
99405 42222
ANNA NAGAR
97909 25554
ADYAR
9790925550

www.nacedu.com

The Art of Guard

SELinux Logging

In the previous article in this series, we looked at allow rules in an SELinux policy. This month we'll discuss SELinux error logs in order to decipher them and take corrective action.

Most users and sysadmins would have come across a situation in which a desired application/program does not behave as expected in a system running SELinux. I have seen that, often, the blame falls squarely on SELinux—

whether it is the culprit or not.

To successfully troubleshoot an issue with application performance in an SELinux-enabled system, follow the steps shown below to diagnose the problem:

1. Determine whether the misbehaviour/error is due to SELinux.
2. If not, troubleshoot the application.
3. If yes, decipher SELinux error logs to understand why the denial occurred and take corrective action.

This article addresses the third point—i.e., how to

decipher SELinux logs and take corrective action.

To determine whether the application misbehaviour/error is due to SELinux, switch SELinux to permissive mode:


```
[root@vbg ~]# setenforce 0
```

Now re-execute the earlier task.

If you still face errors, SELinux is not responsible for the malfunction. The troubleshooting efforts need to be directed elsewhere (the application log files would be a good place to start).

If indeed the application behaved as desired on changing SELinux mode to permissive, it is time that you troubleshoot the SELinux logs.

Any denials by SELinux are recorded in the log files as Access Vector Cache (AVC) denials, since AVC is used by the rules engine.



Host:usr	Message	Date	Source Type	Target Type	Object Class	Permissions	Executable	Command	Other
Denied		Aug 22 15:42:34	unconfined	openoff_bin	file	getattr	firefox-bin	dev=hda6 timestamp=1250669954.374	
Denied		Aug 24 00:13:47	httpd_t	tmp_t	file	getattr	httpd	dev=hda3 timestamp=1250389027.148	
Denied		Aug 24 06:13:47	httpd_t	tmp_t	file	read	httpd	dev=hda3 timestamp=1250389027.165	
Denied		Aug 28 05:47:43	httpd_t	tmp_t	file	getattr	httpd	dev=hda3 timestamp=1250569063.621	
Denied		Aug 28 05:47:43	httpd_t	tmp_t	file	read	httpd	dev=hda3 timestamp=1250569063.622	
Denied		Aug 07 13:25:57	unconfined	openoff_bin	file	read	python	dev=hda6 timestamp=1249631637.161	
Denied		Aug 07 13:25:57	unconfined	openoff_bin	file	getattr	python	dev=hda6 timestamp=1249631637.162	
Denied		Aug 08 14:06:32	unconfined	openoff_bin	file	read	python	dev=hda6 timestamp=1249720592.394	
Denied		Aug 08 14:06:32	unconfined	openoff_bin	file	getattr	python	dev=hda6 timestamp=1249720592.409	
Denied		Aug 08 14:07:28	unconfined	openoff_bin	file	read	python	dev=hda6 timestamp=1249720648.061	
Denied		Aug 08 14:07:28	unconfined	openoff_bin	file	getattr	python	dev=hda6 timestamp=1249720648.062	
Denied		Aug 30 09:34:59	unconfined	openoff_bin	file	read	python	dev=hda6 timestamp=1249677099.054	
Denied		Aug 30 09:34:59	unconfined	openoff_bin	file	getattr	python	dev=hda6 timestamp=1249677099.068	
Denied		Aug 30 16:26:02	unconfined	openoff_bin	file	read	python	dev=hda6 timestamp=1249601762.667	
Denied		Aug 30 16:26:02	unconfined	openoff_bin	file	getattr	python	dev=hda6 timestamp=1249601762.681	
Denied		Aug 30 18:10:22	unconfined	openoff_bin	file	read	python	dev=hda6 timestamp=1249508022.065	
Denied		Aug 30 18:10:22	unconfined	openoff_bin	file	getattr	python	dev=hda6 timestamp=1249508022.065	

Policy Version: v22 (binary, mbs) Log Messages: 7193/7193 Dates: Aug 07 12:45:56 - Aug 22 14:12:33 Monitor status:

Figure 1: The seaudit tool

SELinux Log files

If the Linux Auditing System (the *auditd* daemon) is running, SELinux denials are logged into the audit log file. The default audit log file is */var/log/audit/audit.log*. In a situation where the *auditd* daemon is not running, AVC denials are logged in */var/log/messages*.

In my opinion, it is prudent to enable the audit daemon and maintain a separate log file. The audit RPM should be installed by default on most Red Hat Enterprise systems.

If not installed, do so. In my system running RHEL 5, the RPM is *audit-1.3.1-1.el5*.

Also, do ensure that the *auditd* daemon is enabled to start at boot time. You can do that by issuing the following command:

```
[root@vbg ~]# chkconfig auditd on
```

To ensure that the *auditd* daemon has been enabled at start-up, issue the following command:

```
[root@vbg ~]# chkconfig --list auditd
auditd    0:off 1:on 2:on 3:on 4:on 5:
on 6:off
```

As you can see, *auditd* will be executed whenever my system boots in run levels 3 or 5.

Once *auditd* is running, SELinux

logs are written to the audit log file (generally */var/log/audit/audit.log*).

A typical snippet of an AVC denial message, apart from the time stamp in my log file, is given below:

```
avc: denied { read } for pid=3579
comm="httpd" name="index.html" dev=hda3
ino=33001 scontext=user_u:system_r:httpd_t:s0
tcontext=system_u:object_r:tmp_t:s0 tclass=file
```

Let us look at the log and decipher its meaning:

```
avc : [result of SELinux rules] { access type
} for pid="subject - Process Id Number"
comm="subject - program command"
name="object - name" dev="object - device"
ino="object - inode number" scontext="subject
- Security Context of Source" tcontext="object
- Security Context of Target" tclass=" object
- Class Type of target"
```

The above error can be interpreted by classifying the subject, object details and understanding the operation that the subject was performing on the object.

- Subject details in the log file:
 - Subject Command = httpd
 - Subject PID = 3579
 - Subject Security Context = user_u:system_r:httpd_t:s0
- Object details in the log file:
 - Object Name = index.html
 - Object Class = file
 - Object Partition = hda3
 - Object Inode Number =

RHCA (Red Hat Certified Architect)
Training & Exam
At FOSTERING Linux
By Mr. Arun Eapen, RHCA
(Red Hat Certified Architect)
Sr. Technical Consultant,
Red Hat India

RH 401: Red Hat Enterprise
Deployment, Virtualization and
System Management
 Training Dates: 20-23 Sept09
 Exam: 6 Oct 09

RH 436: Red Hat Enterprise
Clustering and
Storage Management
 Training Dates: 24-27 Sept 09
 Exam: 4 Oct 09

RH442: Red Hat Enterprise
System Monitoring &
Performance Tuning
 Training Dates: 28Sept-1Oct 09
 Exam: 5 Oct 09

New Batches:

RHCE: Call for details

RHCSS: Call for details

Jboss: 12th Sept 09

PHP with LIVE project: 5th Sept 09

Cluster: 12th Set 09

Mail Server: 20th Sept 09

Python: 19th Sept 09

Exam Dates

EX 333: 15th Sept 09

RHCE: 29th Sept 09

EX 429: 29th Sept 09

EX 423: 30th Sept 09

FOSTER Linux Inc.
 Free & Open Source Training Education & Research

SCO M-22, Second Floor Old DLF Colony,
 Sector -14 Gurgaon 122 001,
 Haryana, Tel: +91-124-4268187, 4080880

55, Deepak Building 13, Nehru Place
 New Delhi 110 019 Tel: +91-30330047 / 49
 Email Id: info@fosteringlinux.com

redhat
 TRAINING PARTNER

- 33001
- Object Security Context = system_u:object_r:tmp_t:s0
3. Access details in the log file:
Operation being Performed by Subject{ described in 1) above } on Object{ described in 2) above } = read
 4. Result: Result of the Operation{ described in 3) above } = denied
- The above can also be viewed in an easy-to-understand GUI. To be able to view the log file in GUI mode, install the *setools-gui* package.

Once installed, you can view your log file by executing the *seaudit* command. A GUI window shall open, which is similar to Figure 1.

In the file menu, select *Open Log* and then select */var/log/audit/audit.log*.

Once you open the log file, you can see the details of error logs generated graphically, sorted in various columns. You can right click on the message of your choice and look at the original log message.

Resolving SELinux errors

Once you have decoded the error from the log file, SELinux provides a plethora of tools to help solve the errors. Depending on the solutions and the security level desired by a systems administrator, the possible solutions can be applied to allow the action. The *setroubleshoot* RPM will be required for this.

The *setroubleshoot* RPM is an important tool in the SELinux systems administrator's toolkit. If it is not installed in your system, please go ahead and install it. *sealert* is one of the utilities provided.

To understand *sealert* usage, issue the following command:

```
[root@vbg selinux]# sealert -h
-b --browser      Launch the browser
-h --help         Show this message
-s --service      Start sealert as a dbus
service
-S --noservice    Start sealert without dbus
```

service as stand alone app

```
-l --lookupid id   Lookup alert by id
-a --analyze file  Scan a log file, analyze it's AVC's
-H --html_output   Output in html
-v --verbose       Start in verbose mode
```

Let us first analyse the log file */var/log/audit/audit.log* using *sealert*. To do so, enter the command that follows:

```
[root@vbg selinux]# sealert -a /var/log/audit/audit.log
```

Depending on the number of errors in the log file, an output containing the possible solutions to AVC denial errors in the log file will appear on the screen. To analyse these solutions, you can redirect the output to a file:

```
[root@vbg selinux]# sealert -a /var/log/audit/audit.log > \
/tmp/my-selinux-error-solutions.txt
```

Each error and possible solution is mentioned in the text file (*/tmp/my-selinux-error-solutions.txt*) separated by a line of dashes. The AVC error example discussed in the above section has been analysed by *sealert* as:

Summary

SELinux is preventing the /usr/sbin/httpd from using potentially mislabeled files (/var/www/html/index.html).

Detailed Description

SELinux has denied /usr/sbin/httpd access to potentially mislabeled file(s) (/var/www/html/index.html). This means that SELinux will not allow /usr/sbin/httpd to use these files. It is common for users to edit files in their home directory or tmp directories and then move (mv) them to system directories. The problem is that the files end up with the wrong file context which confined applications are not allowed to access.

Allowing Access

If you want /usr/sbin/httpd to access this files, you need to relabel them using `restorecon -v /var/www/html/index.html`. You might want to relabel the entire directory using `restorecon -R -v /var/www/html`.

Additional Information

```
Source Context      user_u:system_r:httpd_t:s0
Target Context      system_u:object_r:tmp_t:s0
Target Objects      /var/www/html/index.html [ file ]
Affected RPM Packages
Policy RPM
Selinux Enabled
Policy Type
MLS Enabled
Enforcing Mode
Plugin Name         plugins.home_tmp_bad_labels
Host Name
Platform
Alert Count         2
Line Numbers        8796,8797,8798
```

31755,0-1 70%

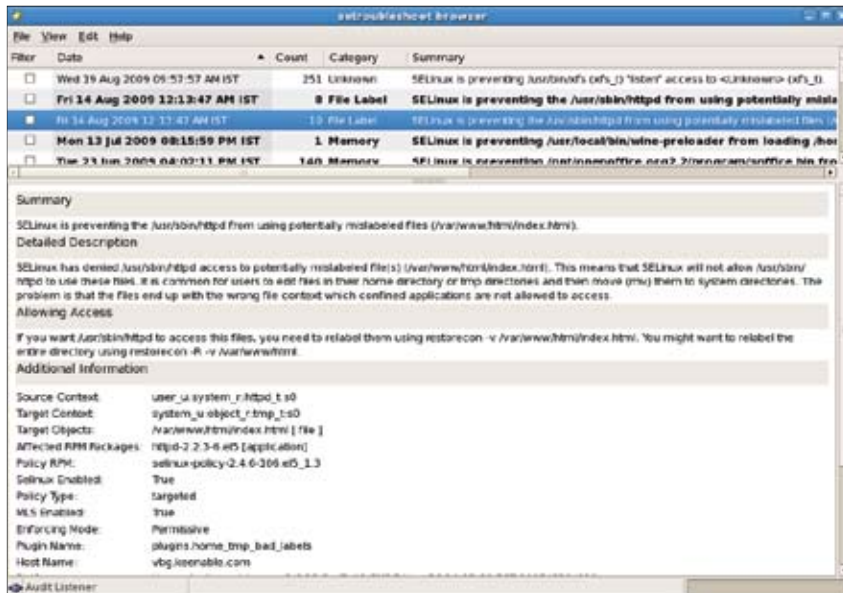


Figure 2: *sealert* can also show errors in a graphical window

Raw Audit Messages

```

avc: denied { getattr } for comm="httpd"
dev=hda3 egid=48 euid=48
exe="/usr/sbin/httpd" exit=0 fsgid=48 fsuid=48
gid=48 items=0 name="index.html"
path="/var/www/html/index.html" pid=3579
scontext=user_u:system_r:httpd_t:s0
sgid=48 subj=user_u:system_r:httpd_t:s0
suid=48 tclass=file
tcontext=system_u:object_r:tmp_t:s0 tty=(none)
uid=48

```

The above output from *sealert* helps a systems administrator to:

1. Look at the raw log message
 2. Understand why the AVC denial occurred (giving detailed description and summary)
 3. View possible solutions to avoid the AVC denial (by allowing access)
 4. View additional information
- Executing the 'Possible Solution' would generally solve the AVC denial error.

If you execute the solution mentioned below:

```
[root@vbg selinux]# restorecon -v /var/www/html/index.html
```

...the security context of this file would change to *system_u:object_r:httpd_sys_content_t:s0*. This would

remove the AVC denial and you would be able to browse your HTML page.

For those of us inclined to use the GUI, *sealert* can be initiated in the GUI mode, by giving the following command:

```
[root@vbg selinux]# sealert -b &
```

Please refer to Figure 2, which shows the output for the same error graphically in the browser.

The *setroubleshoot* RPM also includes an *init* script called *setroubleshoot*. By enabling the *setroubleshoot* script on startup using the *chkconfig* utility, *setroubleshoot* runs as a daemon in the background. Every time an AVC denial occurs, a pop-up appears in the system tray (assuming you are working in the GUI mode). This helps systems administrators to troubleshoot SELinux errors on-the-fly.

At times, you might not be

willing to choose the solution suggested by the *sealert* tool and would rather create your own *allow* rules in the SELinux policy. For example, I would want my Web server *httpd* to be able to read files of type *tmp_t* rather than restrict it to read only files of type *httpd_sys_content_t*.


In such situations, a new *allow* rule will have to be added to the Security Policy. This *allow* rule will permit a subject of *httpd_t* to read object files of type *tmp_t*.

The above would give you great flexibility in creating your own security environment and not just follow the standard policies that came with the system. At other times, you would require to create security policies for other applications that you may have installed on your system, for example, an Oracle database.

Do you need to build the entire SELinux policy for this? That would appear to be an uphill task for most systems administrators. Therefore, SELinux allows admins to create their own modules.

Instead of modifying the core policy, you can build modules of your own that can be loaded on top of the core policy. In these modules, you can declare your own types and rules.

In the next article in the series, we will look at creating SELinux modules, compiling them and loading them.

I would look forward to your feedback on the articles. Please feel free to drop me a mail for suggestions on improving the content and making it more useful. **END** 

By: Varad Gupta

Varad is an open source enthusiast who strongly believes in the open source collaborative model not only for technology but also for business. India's first RHCSS (Red Hat Certified Security Specialist), he has been involved in spreading open source through Keen & Able Computers Pvt Ltd, an open source systems integration company, and FOSTERING Linux, a FOSS training, education and research training centre. The author can be contacted at varad.gupta@fosteringlinux.com

Sneak Into Your Thumb Drive from the Cloud



Let us explore cloud computing and learn how to set up the cloud that can virtualise your USB device.

Cloud computing is the new paradigm of the technical world, allowing companies to reduce infrastructure costs and also part of software licence fees. This computing model is not just for the enterprise world, but offers innovative solutions to the end users as well.

For most of us, computing has become a part of our daily activities. Typically, the computer is used for social networking, maintaining financial records, or even playing games. This computing comes at a cost that covers procurement, maintenance and various other aspects that do tend to hassle users. The evolution of cloud computing—which means computing as a

service—is a viable solution that addresses these problems.

Cloud computing has borrowed its basics from several other computing areas and systems engineering concepts. The most common cloud computing servicing models are software as a service (SaaS), the platform as a service, and infrastructure as a service – all of which are provided to users over the Internet. Figure 1 provides a better understanding of all the services that cloud computing could provide.

We can use any of these services from our desktops, mobile phones or thin clients (see Figure 2). The service providers will be remote (or ‘in a cloud’) with respect to the user.

Now let's look at how a device like a USB thumb drive can be used in a cloud set-up.

Peripheral virtualisation

In the course of using a computer today, we connect devices like a thumb drive, printer, Web camera, etc, besides a mouse and keyboard, to a PC. But in a cloud computing set-up, when the computing service is in a remote location, how does an end user use such peripherals? In an ideal cloud set-up, these devices get



virtually connected to a cloud server even though the device is connected to the client, which is at the user's home set-up. These ideal cloud servers and home set-ups are available in Linux (and also other operating systems).

As a Linux hacker, you should be aware of what is happening inside the kernel driver layer that enables virtual peripheral connectivity. To understand this, let's look at how to modify the Linux kernel and virtually connect USB devices to a cloud server. One important aspect to take care of during virtualisation is that timing requirements have to be satisfied as per the protocol specification, besides transparency to the user that it's a virtual device. Thus, when virtualising a USB device, care has to be taken to ensure the quality of service and that the user doesn't feel like the device is remotely connected.

To virtualise a USB device in a remote cloud server, 'information' that helps virtualise the device has to be made available over the IP network. The current driver design has to be modified to suit this requirement. To understand this, let us first analyse the USB driver architecture and then understand where to hook on this information to the IP layer.

A Linux USB host driver consists of three major layers. The first low-level module is the Host Controller Driver (HCD), which implements hardware specific routines. The next portion is the USB Core Driver implementing USB-specific functions. The final layer is the application layer. Figure 3 illustrates the USB driver architecture.

When our client desktop connects to a cloud server over the Internet, the device information and the data from the USB device has to be transferred through the network to a server on the remote cloud. This device information and data has to be tapped from the appropriate layer so that the virtual device can function properly. Figure 4 shows how the information could be tapped in a USB device.

USB virtualisation in your home lab

Having looked at all internal kernel changes, let us now explore how to set up a cloud-like environment in your home lab. To do the set-up, prepare two Linux machines and a USB peripheral like a thumb drive.

The open source ecosystem is rich with projects that allow us to explore multiple ideas and this applies to the current requirement too. The USB/IP project will help you to understand a possible approach of how a USB device can be virtualised across a network. The USB/IP ideally shares USB devices between computers by encapsulating USB transfers into TCP/IP payloads and transmits them between computers.

Here is a quick overview of how to set up a USB/IP at home. [For detailed steps, refer the *README* file or hack the source code.] In a USB/IP, the PC that exports a virtual device is referred to as a server and the PC that uses a virtual device, as the client.

The first step in the process is to download the source packages from <http://usbip.sourceforge.net> in both the computers. Compile the source in the driver folder and install the server drivers following the steps given in the *README* file [<http://usbip.svn.sourceforge.net/viewvc/>

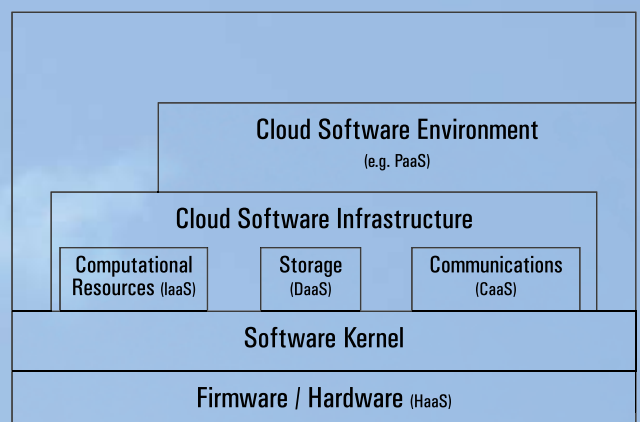


Figure 1: Five layers of cloud computing




Figure 2: Possible cloud setup

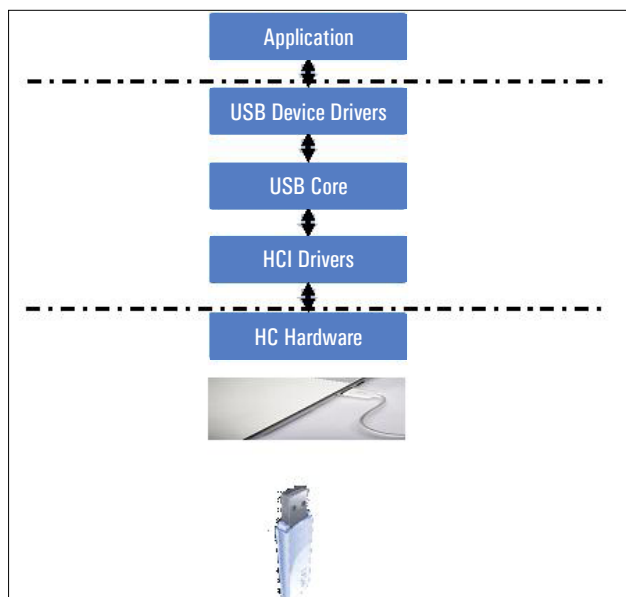


Figure 3: USB Linux software architecture in a normal setup

usbip/linux/trunk/drivers/README]. Compile the USB/IP utils from the *src/* directory and install the USB/IP daemon from the *src/* folder as detailed in the following *README* file: <http://usbip.svn.sourceforge.net/viewvc/>





Figure 4: Possible USB software architecture in a cloud setup

usbip/linux/trunk/src/README. Follow the same process in the client PC as well.

After the successful installation of drivers, insert the thumb drive in the server PC. Go through the steps in the following *README* file to export the thumb drive and access the client PC: <http://usbip.svn.sourceforge.net/viewvc/usbip/linux/trunk/src/README>.

The USB/IP is an experimental project that was developed to extend the peripheral bus. Enthusiastic readers should refer to the following URL to get more insights into the motivation behind the project and the achievements so far: <http://www.usenix.org/events/usenix05/tech/freenix/hirofuchi/hirofuchi.html/index.html>.

As new markets emerge with new technologies and different business models, traditional software architectures need to be looked at again. New innovative software architectures like USB/IP have to evolve to suit emerging technologies and business models, providing the user a better experience in terms of features and usability. **END** 

By: Rajaram R

The author welcomes your comments and feedback at rera_raja@yahoo.com

Python for Research

An Initiation

Let's get started with NumPy, which is supposed to be "the fundamental package needed for scientific computing with Python."


The purpose of research is to try and explain something to the world based on certain observations. The challenge is to come up with an explanation before anyone else does. You, as a researcher, are confronted with a volume of data. You need to model it and demonstrate that your model is correct.

How likely are you to hit upon a valid model on your first try? Not very, unless the problem is trivial. So, the challenge is not just to prove that the correct model works. The even greater challenge is to find out as quickly as possible, whether your model is the correct solution, and if not, to eliminate it and move on. This is precisely where Python is superb.

Well, there are some issues.

Python works with dynamic types. A list can contain any mixture of data types. The research problem deals with data of one type. The mixture of data types, obviously, allows for the data to be of the same type as well. However, there is a price. The type of data has to be checked at run time. As a consequence, for massive computation, Python can be virtually unusable.

So, Jim Hugunin, a graduate student at MIT, created Numeric in 1995. He then went on to create Jython and IronPython. Numeric continued to grow in the scientific and engineering community. Many students and scientists built their modules around the Numeric module. The project's name on SourceForge.net was NumPy.



This work has now become a part of the much broader SciPy project. You can find out more about it at www.scipy.org/History_of_SciPy.

Python does not come with the NumPy module included by default. It is, however, available in various distributions. You will need to make sure that it is installed. Familiarity with core Python is assumed.

NumPy one-dimensional arrays or vectors

The best way to learn is to try a little code. Suppose you need to multiply corresponding elements of two arrays. How would the code and performance be different? Try the following small benchmark:

```
N=1000000
a=[i for i in range(N)]
b=[0.5*i for i in range(N)]
for t in range(10):
    ab = [a[i]*b[i] for i in range(N)]
```

You have defined two arrays of a million elements each. Your code creates a new array, which is the product of the corresponding elements. You repeat it 10 times so that the execution times can be compared reasonably.

Now, you can write the same code using NumPy arrays.

```
import numpy as np
N=1000000
a=np.array([i for i in range(N)])
b=np.array([0.5*i for i in range(N)])
for t in range(10):
    ab = a*b
```

The method `array` in the NumPy module converts a Python list into an array object of uniform type elements. There are a number of operations defined on these objects—for example, multiplication results in a new array, which is the product of the corresponding elements.

On my system, the core Python program took 8.9 seconds while the NumPy version took 1.7 seconds.

Now, try the following code:

```
>>> a = [1,2,3]
>>> b = [3,2,1]
>>> a + b
```

Before you give your answer, try the following as well:

```
>>> s = ['a','b','c']
>>> a + s
```

The only reasonable, consistent answer is concatenation of the two lists. Now, try the same thing with NumPy arrays:

```
>>> a=np.array([1,2,3])
>>> b=np.array([3,2,1])
```

```
>>> s=np.array(['a', 'b', 'c'])
>>> a + b
array([4, 4, 4])
>>> a + s
```

Since the elements in the array are of the same type, summation of *a* and *b* is very reasonable. Now, what should happen in the case of adding *s* to *a*? Understandably, it should raise an exception, as the addition of a number to a string is not meaningful.

You need to compute the sine of a list of numbers. This can also be a good test to benchmark. So, try the following code:

```
import math
N=1000000
a=[i*2*math.pi/N for i in range(N)]
for t in range(10):
    sina = [math.sin(v) for v in a]
```

You have defined an array of one million equally spaced values between 0 and 2π . As before, create a new array containing the sine of the corresponding value in the first array. Repeat it 10 times so that the run time is large enough.

The corresponding program using NumPy will be:

```
import math
import numpy as np
N=1000000
a=np.array([i*2*math.pi/N for i in range(N)])
for t in range(10):
    sina = np.sin(a)
```

As before, you have converted a Python list into a NumPy array. You can pass an array to a function and its meaning is very clear. It creates a new array whose values are the function of each of the corresponding elements of the array, passed as a parameter.

In this case, on my system, the core Python version takes 7.3 seconds while the NumPy version takes 1.8 seconds.

NumPy multi-dimensional arrays and matrices

A list of lists would be a two dimensional array. Try another little test:

```
r = 1280*[1]
m = 1024*[r]
for t in range(10):
    m2 = []
    for r in m:
        m2.append([.5*e for e in r])
```

Define a list of 1,024 rows. Each row has 1,280 elements, with a value of 1. The task is to scale all

elements with a constant.

The corresponding program using NumPy would be:

```
import numpy as np
r = 1280*[1]
m = np.array(1024*[r])
print m.shape
for t in range(10):
    m2 = .5*m
```

As before, the array method converts a Python list into an array. Printing the shape of the array confirms that it is a two dimensional array of size 1024x1280. The core Python program takes 5.9 seconds while the NumPy application takes 0.7 seconds.

But notice how easy it is to multiply all elements in an array with a constant. You can use NumPy to define the matrix with all elements being 1, as follows:

```
import numpy as np
m = np.ones((1024,1280))
print m.shape
for t in range(10):
    mp = .5*m
```

This version runs in 0.4 seconds.

It is reasonable to assume that the addition of two arrays of the same shape is possible and the corresponding elements will be added. Explore the following:

```
>>> m = np.array([[1,2,3], [ 4,5,6]])
>>> n = np.array([[1,2,3], [ 4,5,6]])
>>> m.shape
(2, 3)
>>> n.shape
(2, 3)
>>> m + n
array([[ 2,  4,  6],
       [ 8, 10, 12]])
>>> n.shape = (6,1)
>>> m + n
```

You are creating two arrays of size 2x3 and adding them. The result is as expected. NumPy allows an array to be reinterpreted by changing the shape of the array. In

the example above, the array *n* is reinterpreted as a one-dimensional array. Now, if you try to add *m* and *n*, it should fail and it does.


Multiplication of two-dimensional arrays could be interpreted as a matrix multiplication. However, that would not be consistent across various dimensions. Hence, in NumPy, multiplication of two arrays is allowed if their shapes are identical and correspond to the creation of a new array whose elements are a product of the corresponding elements of the two arrays. You can try the following:

```
>>> m = np.array([[1,2],[3,4]])
>>> n = np.array([[1,2],[3,4]])
>>> m*n
array([[ 1,  4],
       [ 9, 16]])
```

If you want to work with matrices, you need to define a two-dimensional array as a matrix. Try the following example—compare and explore:

```
>>> m = np.matrix([[1,2],[3,4]])
>>> n = np.matrix([[1,2],[3,4]])
>>> m*n
matrix([[ 7, 10],
        [15, 22]])
>>> m.shape=4,1
>>> n.shape = 4,1
>>> m*n
>>> n*m
```

Use the built-in help to know more about the possibilities available for arrays, matrices and a lot more.

It should be clear that if an operation can be applied to an array, the performance gains over core Python could be substantial. NumPy provides a lot of functionality, which can be very useful when analysing large volumes of data, giving researchers an edge to beat the competition! 

By: Dr. Anil Seth

The author is a consultant by profession and can be reached at seth.anil@gmail.com

TechnoMail - Enterprise Email Server
Anti SPAM, Anti Virus, Email Content Filtering

Firewall, Internet Access Control
Content Filtering, Site Blocking

Bandwidth Management System

Managed Email Hosting Solutions

TechnoInfotech®

1, Vikas Permisses, 11 Bank Street, Fort, Mumbai, India - 400 001. Tel.: 91-22-6633 8900 Ext. 324. info@technoinfotech.com



Tips & Tricks



Count the number of SSH commands

The following command will find the number of the SSH connections from a local to a remote machine:

```
ss | grep -vE 'State' | awk '{print $1 "--> " "Local("$4)" " to " "Remote("$5")}' | grep -e ssh | wc -l
```

The following command will print out the SSH connections from the local to the remote machine:

```
ss | grep -vE 'State' | awk '{print $1 "--> " "Local("$4)" " to " "Remote("$5")}' | grep -e ssh
```

—Shreyas Parikh, shreyu82@gmail.com



Investigation of sockets

The following command will display the names of the ports that are connected to a remote server:

```
ss | grep -vE 'State' | awk '{print $5}' | cut -d ':' -f2 | sort | uniq
```

The command will print only the established port names.

—Shreyas Parikh, shreyu82@gmail.com



How to record a Linux session

If you want to record your Linux session with the commands given and their outputs, use the script command as follows:

```
root@localhost#script <filename.log>
```

It basically forks a new shell so that if you want to stop a session recording, just type exit or press Ctrl+d. You can find the recorded session in the ~/session.log file.

Another interesting use of the script utility is redirecting your typed characters to another user's

terminal — someone who is logged on to Telnet/SSH. To do this, you must learn the devpts number of that person (which can be obtained from the output of the who command). And if it is 7, for example, then the following command...

```
root@localhost#script /dev/pts/7
```

...will redirect all typed characters to that user's terminal window.

—Selvaraj, selvaraj2005@gmail.com



Skip or bypass a fsck

Generally, fsck is run automatically at boot time when the system detects that a filesystem is in an inconsistent state, indicating a non-graceful shutdown, such as a crash or power loss.

We can bypass it on the next reboot by the following three methods:

- 1) Using the shutdown command

```
# shutdown -rf now (here -f option will skip the fsck.)
```

```
# shutdown -rF now (here -F option will force the fsck on next reboot)
```

- 2) By editing the kernel parameter in /boot/grub/grub.conf (or menu.lst). Open grub.conf or menu.lst and find the kernel line. Put fastboot at the end of this line. An example kernel line is shown below:

```
kernel /vmlinuz-2.6.18-92.1.22.el5 ro root=LABEL=/ console=ttyO console=ttyS1,19200n8 fastboot
```

- 3) By updating /etc/fstab file. You will see two numbers at the end of the line for each partition. Change the second number to a 0 (zero digit). This will have the system mount the partition, but will not run a check when booting. If the sixth field is neither present

nor a zero, a value of zero is returned and fsck will assume that the filesystem does not need to be checked. For example:

```
LABEL=/disk3 /disk3 ext3 defaults 0 0
```

—*Jasvendarsingh M. Chokdayat,*
theindianjash@gmail.com



Temporary IP address and gateway

If you need to quickly text a network (IP) address, here's how you can do it:

```
ifconfig eth0 192.168.1.100 netmask 255.255.255.0
```

To add a gateway:

```
route add default gw 192.168.1.1
```

The above commands will disappear the moment you reboot your system.

—*Melvin Lobo, melvin.lobo@axonnetworks.com*



Search in man pages

To search for a word in the description of the manual pages of all Linux commands installed, you can use the *apropos* command. Syntax: *apropos <string/word>*

For example:

```
$ apropos list
acl (5) - Access Control Lists
appres (1) - list X application resource database
chacl (1) - change the access control list of a file or directory
column (1) - columnate lists
dir (1) - list directory contents
fc-list (1) - list available fonts
fpit (4) - Fujitsu Stylistic input driver
fslsfonts (1) - list fonts served by X font server
getfacl (1) - get file access control lists
<<output truncated>>
```

As you can see, it lists all the commands that have the word 'list' under their description section in their corresponding man pages.

—*K. Indraveni, indraveni.chebolu@gmail.com*



Fun with the terminal

Paste the following command on a terminal and press *Enter*:

```
yes "$(seq 232 255;seq 254 -1 233)" | while read i; do printf "\x1b[48;5;${i}m\
```

```
n";sleep .01;done
```

Do remember to reset your terminal window after exiting the loop.

—*Shreyas Parikh, shreyu82@gmail.com*



Prevent Linux from remembering your *sudo* password

In order to do this, use the following code:

```
sudo visudo
```

..and add this line to the file:

```
Defaults timestamp_timeout = 0
```

You may change 0 to any number representing the minutes you may want your password to be 'remembered', or let it be 0, in which case you will need to type your password each time you type *sudo*

—*Remin Raphael, remin13@gmail.com*



Prevent users from changing their password

The */usr/bin/passwd* file usually has the SUID permission set:

```
# ls -l /usr/bin/passwd
-r-s--x--x 1 root root 19348 Sep 7 2004 /usr/bin/passwd
```

As you can see, it's 'r-s' instead of 'r-x' for the file owner, i.e., the root. SUID files get the (limited) permissions of the file owner instead of the user who runs it. So we need to remove the SUID for that command, as follows:

```
chmod u-s /usr/bin/passwd
```

Now, normal users won't be able to change their own passwords—and only the root user will be able to do it for them.

—*Govindarajalu, govind.rajalu@gmail.com*



Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in LFY so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at www.linuxforu.com. The sender of each published tip will get an LFY T-shirt.



Typo Bugs

Can typing mistakes (typos) cause bugs? Yes, they can! We'll look at some common C programming mistakes in this column.

What does the following program print?

```
int val = 2;
switch(val) {
case 1:
    printf("case 1");
    break;
defalut:
    printf("in default");
}
```

Here, you would expect "in default" to get printed; but it does not print anything and the default case does not execute. Why? The keyword "default" was mistyped as "defalut". Why does the compiler still compile the code without complaining? In C (as in many other languages), a label (as a target of the *goto* statement) is a name followed by a : (colon). In this case, "defalut:" was treated as a label and hence the bug!

In the May 2009 JoP column, we discussed a typographical (typing) mistake that resulted in the loss of the Mariner-I rocket because of mistyping a . (full stop) instead of a , (comma). So, we'll go further and discuss some common typing mistakes in C that we as programmers make while writing code.

- 1) Typing 'l' (lowercase 'L') instead of 1 as in the statement:

```
long val = 9l;
```

In this case, the constant is 9 (nine) and not 91 (ninety-one) since 'l' is a suffix to indicate a long variable.

- 2) Typing = instead of == , which makes a comparison expression into an assignment statement, as in:

```
if( a = 0 ) // the programmer wanted to check if(a == 0)
```

This mistake is so very common in C that a "defensive programming" technique is to reverse the condition, as in, `if(0 == a)`, so that compiler will catch it if we make a mistake.

- 3) Typing == instead of = , which makes an assignment statement into a comparison statement, as in:

```
x == 2;
```

In fact, this example is from Peter van der Linden's book *Deep C Secrets* in which he talks about a bug that was holding up a \$20 million sales deal at Sun. It turned out to be this mistake in the code in the I/O library.

- 4) Typing a ; (semi-colon) after a *for* loop by mistake, as in:

```
for(i = 0; i < 10; i++);
printf("hello");
```

And this code prints "hello" once and not ten times.

- 5) Forgetting to type a comma between two strings, which leads to "stringization" of two strings, as in:

```
char *str[] = { "rain" "deer", 0 };
int i = 0;
while(str[i])
    puts(str[i++]);
```

This prints "raindeer" instead of "rain" and "deer".

- 6) Can you find out what's wrong with this code?

```
struct Point {
    int x, y;
}
foo() {}
```

Here, we forgot to put a semicolon after the *struct* definition point. In old C, if we don't provide a return type, it is considered to return *int* by default. However, the definition for function *foo* is that it returns the struct *Point*! (Yes, we can define a struct/class "on-the-fly" as a return type in C/C++!). This typo bug has caused many sleepless debugging nights for programmers worldwide.

- 7) As of this writing, the latest news is about "Tiny typo blamed for massive IE security failure"[1]. The following code and the explanation is from [2]:


```
hr = pStream->Read((void*)&pbArray, (ULONG)cbSize, NULL);
```

I'll give you one more clue – it's a one-character typo. ...The first argument is incorrect. It should be:

```
hr = pStream->Read((void*)pbArray, (ULONG)cbSize, NULL);
```

The extra '&' character in the vulnerable code causes the code to write potentially untrusted data, of size `cbSize`, to the address of the pointer to the array, `pbArray`, rather than write the data into the array, and the pointer is on the stack. This is a stack-based buffer overrun vulnerability.

- [1] http://www.theregister.co.uk/2009/07/30/typo_caused_massive_ms_bug

- [2] <http://blogs.msdn.com/sdl/archive/2009/07/28/atl-ms09-035-and-the-sdl.aspx> 

About the author:

S G Ganesh is a research engineer in Siemens (Corporate Technology). His latest book is "60 Tips on Object Oriented Programming", published by Tata McGraw-Hill. You can reach him at sgganesh@gmail.com.



Sandya Mannarswamy

Welcome to another instalment of CodeSport. In this month's column, we continue our discussion on the false sharing issue in multi-threaded applications. We'll also discuss priority inversion and possible techniques to deal with it.

Thanks to all the readers who responded to the problems we discussed in last month's column. We had given a small code snippet of multi-threaded code, which exhibits false sharing of data. We had asked readers to rewrite the code to avoid the false sharing issue. Congratulations to Nilesh Govande, Arivendu Bhardwaj, Vivek Goel and Miranda Walters for getting the correct answer. Here is the original buggy code snippet from the takeaway problem:

```
int num_tickets_per_thread [N]; //N is the number of
threads in your application

incrementTicketCounter()
{
    int thread_id = pthread_self();
    num_tickets_per_thread[thread_id] ++;
}

void PrintTotalTickets()
{
    //only the main thread calls this function, so no
locking needed
    for (int i = 0; i < N; i++)
        num_total_tickets += Num_tickets_per_thread[i];
    printf ("%d", num_total_tickets);
}
```

Assuming a cache line of size 64 bytes, we see that 16 elements of the array *num_tickets_per_thread* can fit in the same cache line. Recall that a processor updating the array element *num_tickets*

per_thread[thread_id] ends up invalidating the other processor's cache copies because of their co-location in the same cache line. Unlike true sharing where the same datum is accessed by multiple threads, in this case, different threads access different data items. Because of their co-location in the same cache line, the line gets migrated across the different processor caches. This phenomenon is known as false sharing. This results in an increased number of cache misses and will show up as an increase in the execution time of the application.

```
#define CACHE_LINE_SIZE 64 //we assume a 64 byte
cache line size
struct ticket
{
    int number;
    int dummy_pad[CACHE_LINE_SIZE - sizeof(int)];
};

struct ticket num_tickets_per_thread [N];
//N is the number of threads in your application

incrementTicketCounter()
{
    int thread_id = pthread_self();
    num_tickets_per_thread[thread_id].number ++;
}

void PrintTotalTickets()
{
    //only the main thread calls this function, so no
locking needed
    for (int i = 0; i < N; i++)
```

```

    num_total_tickets += num_tickets_per_thread[i].number;
    printf ("%d", num_total_tickets);
}

```

As our readers correctly pointed out, in order to avoid false sharing, each element of the array `num_tickets_per_thread` needs to be placed on a separate cache line. The modified code given below does not exhibit any false sharing. We add a pad of size 60 bytes (assuming a cache line size of 64 bytes) around each element of the array so that each element is located on a different cache line. Hence, updating of one element of the array does not result in invalidating copies of any other element present in caches of other processors.

Issues to be aware of in lock-based programming

We'll now look at some of the issues associated with multi-threaded programming that arise due to the use of locks and how they can be overcome. Locks are used to protect the access to shared resources in a multi-threaded application. While locks are needed to ensure a consistent application state, they also bring forth certain issues. We have already discussed a couple of these, namely deadlock and data races. In today's column, we will look at another potential problem associated with locks, called *priority inversion*.

Consider an application that has two threads -- a higher priority thread T1 and a lower priority thread T2. Both T1 and T2 need to update a shared global counter that counts the number of operations performed in the application. This shared global counter is protected by a lock, L1. Now when the higher priority thread T1 is ready to run, it ends up waiting for the lower priority thread T2, since T2 currently holds the lock L1. Thus even if T1 has a higher priority than T2, T1 ends up waiting for T2 to finish updating the counter and hence release the lock L1, before it can run. Such a situation is known as priority inversion since T1 has a higher priority than T2, but ends up waiting for T2 clearly resulting in a priority inversion. In an application with a large number of threads, this situation can result in the higher priority thread waiting for a long time on lower priority threads because of access to common shared resources.

A famous case of priority inversion

In many cases, priority inversion happens without the programmer being aware of it and it does not have a major impact on systems performance. But there have been cases where priority inversion has caused poor response on performance, especially in realtime systems. A famous case of priority inversion occurred in the software used in the Mars Pathfinder mission in July 1997. The Pathfinder mission took high-resolution colour pictures of the Martian surface and relayed them back to Earth. When the software was employed on

Mars, the programmers back on earth were puzzled by the number of software resets that occurred frequently, leading to poor systems performance. On investigation, it was found that the issue was due to priority inversion.

In the Martian spacecraft, various devices communicated over a data bus. Activity on this bus was managed by a pair of high-priority tasking threads -- BM1 and BM2. One of the bus manager tasks that BM1 communicated through a pipe was a low-priority meteorological science task, MS1. The communication pipe was protected by a mutex L1. BM1 and MS1 needed to acquire the mutex L1 before they could send data over the pipe.

The sequence of events leading to each reset began when the low-priority task MS1 was pre-empted by a couple of medium-priority tasks while it held the mutex L1. While the low-priority task MS1 was pre-empted, the high-priority bus distribution manager thread, BM1, tried to send more data to it over the same pipe. Because the mutex L1 was still held by MS1, the bus distribution manager was made to wait. Shortly thereafter, the other bus scheduler thread BM2 became active. It noticed that the distribution manager BM1 hadn't completed its work for that bus cycle and forced a system reset. VxWorks was the RTOS used on Pathfinder and using the priority inversion workaround VxWorks had, NASA engineers were able to solve the issue remotely. More details on the priority inversion issue encountered in the Pathfinder mission, can be found in David Wilner's keynote talk at the 1997 IEEE real-time symposium: http://www.research.microsoft.com/research/os/mbj/Mars_Pathfinder/Mars_Pathfinder.html

Dealing with priority inversion

One way of dealing with priority inversion is to use a technique known as priority inheritance. With priority inheritance being enforced by the operating system, a lower priority task will inherit the task of any other higher priority task that is waiting on a shared resource currently owned by the lower priority task. Consider a simple example of two tasks T1 and T2 with priorities P1 and P2, respectively, with P1 being higher than P2. Hence T1 is the higher priority task compared to T2. Both access a shared resource protected by a mutex L1.

Assume that T2 holds the mutex L1 currently. Now, if T1 becomes ready and needs to wait for mutex L1, then T2's priority is boosted to P1. This is done so that T2 can finish quickly and release the mutex so that the higher priority task T1 can then run. So T2 temporarily gets a higher priority for a short time. Since T2 inherits the priority from T1, this technique is known as priority inheritance. Priority change due to priority inheritance takes place as soon as the high-priority task begins to be pending; it ends when the resource is released by the lower priority thread.

Another technique to deal with priority inversion is known as priority ceiling. Priority ceiling associates a priority with each resource; the scheduler then transfers that priority to any task that accesses the resource. The priority assigned to the resource is the priority of its highest-priority user, plus one. Once a task finishes with the resource, its priority returns to normal. Both priority inversion and priority ceiling require operating system support. Linux supports priority ceiling via PTHREAD_PRIO_PROTECT protocol and priority inheritance via PTHREAD_PRIO_INHERIT protocol. For more details on techniques to avoid priority inversion, read the following paper:

<http://www-md.e-technik.uni-rostock.de/ma/gol/rtsys/articulos/YCS159.pdf>

Takeaway problem for this month

This month's takeaway problem comes from Nilesh Govande. Thank you Nilesh, for sending in the question. Given a two-dimensional $N \times N$ array of positive and negative integers, find the sub-rectangle with the largest sum. The sum of a rectangle is the sum of all the elements in that rectangle. In this problem, the sub-rectangle with the largest sum is referred to as the *maximal sub-rectangle*. A sub-rectangle is any contiguous sub-array of size 1×1 or greater, located within the whole array. As an example, the maximal


sub-rectangle of the array:

0	-2	-7	0
9	2	-6	2
-4	1	-4	1
-1	8	0	-2

...is in the lower-left-hand corner:

9	2
-4	1
-1	8

...and has the sum of 15.

If you have any favourite programming puzzles that you would like to discuss on this forum, please send them to me. Feel free to send your solutions and feedback to sandyasm_AT_yahoo_DOT_com. Till we meet again next month, happy programming! 

About the author:

Sandya Mannarswamy. The author is a specialist in compiler optimisation and works at Hewlett-Packard India. She has a number of publications and patents to her credit, and her areas of interest include virtualisation technologies and software development tools.



Calling...

Share your knowledge

Inviting FOSS experts to write articles on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of www.openITis.com (aka linuxforu.com) we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at

lfyedit@efyindia.com

Few topics that top our list:

- ◆ Tips 'n' Tricks for software developers or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools

A Voyage to the Kernel



Part 16

Segment: 3.5, Day 15

We are dedicating today's voyage to learning more theory and clarifying the doubts of beginners.

Inter-Process Communication (IPC)

The Linux Inter-Process Communication (IPC) mechanism is an essential tool that allows the 'sync operation' between processes and facilitates sharing of resources. IPC is also employed for exchanging information with another program. Thus, this is implemented using shared resources, kernel data structures and wait queues. Refer to Figure 1 for the subsystem structure of IPC.

Here are the main IPC implementation methods (note: the following are generalised descriptions):

- **Signals:** These are the oldest form of UNIX IPC and they are actually asynchronous messages sent to a process.
- **Wait queues:** The system can use this to put a process in sleep mode, if the corresponding operation is not yet completed (for example, bottom-half handling by the process scheduler).
- **Pipes (and named pipes):** By using a pipe connection, a connection-oriented bi-directional data transfer between any two processes (or, via a named pipe in the filesystem) can be done.
- **File locks:** This IPC allows a process to declare a file (or part) as read-only to all other processes. Hence the one that holds the lock can only modify it.
- **Unix Domain sockets:** This also is a connection-oriented data-transfer mechanism like the pipe, and the implementation is akin to the INET sockets.
- **System V IPC:**
 - **Semaphores:** This IPC model allows the creation of arrays of semaphores.
 - **Message queues:** This is a connectionless data-transfer model. A message is essentially a sequence of bytes. And by reading the message queues, it is retrieved. This also uses an associated type that restricts the message read.

- **Shared memory:** As the name suggests, this mechanism facilitates accessing a given portion of physical memory by several processes.

The following code shows the notification mechanism for IPC namespaces:

```
#include <linux/msg.h>
#include <linux/rcupdate.h>
#include <linux/notifier.h>
#include <linux/nsproxy.h>
#include <linux/ipc_namespace.h>

#include "util.h"

static BLOCKING_NOTIFIER_HEAD(ipcns_chain);

static int ipcns_callback(struct notifier_block *self,
                          unsigned long action, void *arg)
{
    struct ipc_namespace *ns;

    switch (action) {
        case IPCNS_MEMCHANGED: /* amount of lowmem has
                                changed */
        case IPCNS_CREATED:
        case IPCNS_REMOVED:
            /*
             * It's time to recompute msgmni
             */
            ns = container_of(self, struct ipc_namespace, ipcns_nb);
            /*
             * No need to get a reference on the ns: the 1st job of
             * free_ipc_ns() is to unregister the callback routine
             * blocking_notifier_chain_unregister takes the wr lock to do
             * it.
             * When this callback routine is called the rd lock is held by
             * blocking_notifier_call_chain.
             * So the ipc ns cannot be freed while we are here.
            */
        }
```

```

    */
    recompute_msgmni(ns);
    break;
default:
    break;
}

return NOTIFY_OK;
}

int register_ipcns_notifier(struct ipc_namespace *ns)
{
    int rc;

    memset(&ns->ipcns_nb, 0, sizeof(ns->ipcns_nb));
    ns->ipcns_nb.notifier_call = ipcns_callback;
    ns->ipcns_nb.priority = IPCNS_CALLBACK_PRI;
    rc = blocking_notifier_chain_register(&ipcns_chain, &ns->ipcns_nb);
    if (!rc)
        ns->auto_msgmni = 1;
    return rc;
}

int cond_register_ipcns_notifier(struct ipc_namespace *ns)
{
    int rc;

    memset(&ns->ipcns_nb, 0, sizeof(ns->ipcns_nb));
    ns->ipcns_nb.notifier_call = ipcns_callback;
    ns->ipcns_nb.priority = IPCNS_CALLBACK_PRI;
    rc = blocking_notifier_chain_cond_register(&ipcns_chain,
                                              &ns->ipcns_nb);
    if (!rc)
        ns->auto_msgmni = 1;
    return rc;
}

void unregister_ipcns_notifier(struct ipc_namespace *ns)
{
    blocking_notifier_chain_unregister(&ipcns_chain, &ns->ipcns_nb);
    ns->auto_msgmni = 0;
}

int ipcns_notify(unsigned long val)
{
    return blocking_notifier_call_chain(&ipcns_chain, val, NULL);
}

```

Though we have discussed IPC mechanisms earlier, I think it is necessary to go over some features of these IPC methods. We mainly use a signal to notify a process and it changes the state of the receiving process.

Theoretically, the machine can send these signals to any executing process. The point to note is that if it is a user process, it can send a signal to a process that carries an associated PID (process ID) or GID (in the case of a process

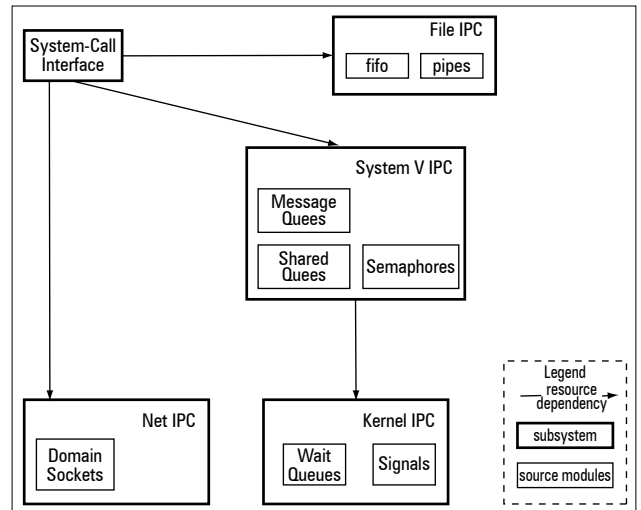


Figure 1: Subsystem structure of IPC




Figure 2: Sub-system dependencies of IPC

group). Before a process initiates, the scheduler will look for a signal and if it finds one, the scheduler uses the *do_signal()* function to handle the signal sent. The second type deals with any process that is in a 'waiting state' (awaiting a kernel event, say the conclusion of a DMA transfer). A process can go for this by just calling the *sleep_on()* function or the *interruptable_sleep_on()* function. Likewise, the *wake_up()* function or *wake_up_interruptable()* function is used to *unlist* from the queue.

In the case of pipes, a file descriptor will refer to the pipe, and one page of memory (a circular buffer) is allotted with the opened pipe. Here, I should clarify that if it tries to read more data than what is available, it will result in a block. Restricting the access to a file is quite important in a typical Linux kernel and this can be done with the help of file-locks. The implementation of UNIX domain sockets is akin to pipes (using the circular buffer). But UNIX domain sockets can offer a separate buffer for each communication direction.

When it comes to semaphores, I must point out that it is, in fact, implemented using wait queues and follows a classical semaphore model, as I mentioned before. Every such thing will have an associated value. *Up()* and *down()* operations can be done using this. It operates in such a way that when its value is zero, the corresponding process (that does the decrement on it) is blocked on the wait queue.

A message queue can be viewed as a linear linked-list, where a process can read/write information (as series of bytes). There are two wait queues in this case. The first wait queue is

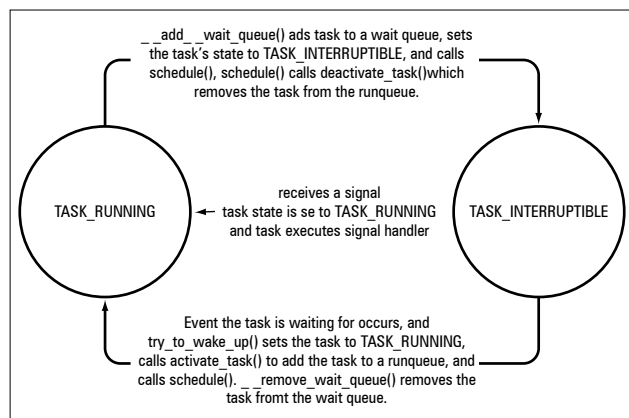


Figure 3: The interrupt process

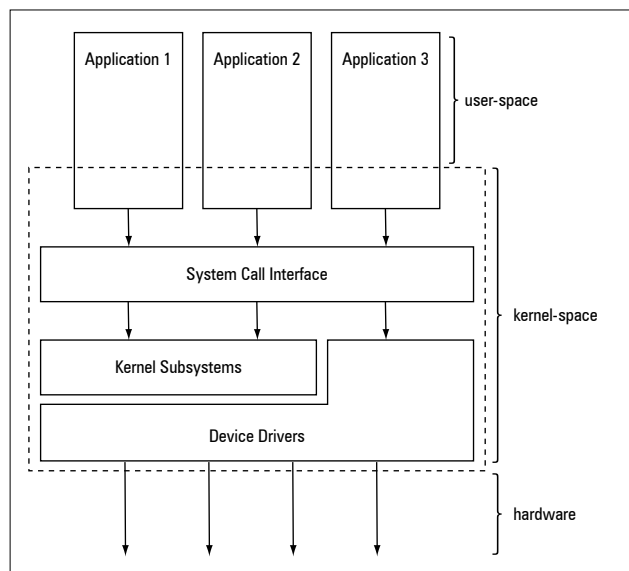


Figure 4: The generalised architecture

for any process that is writing to a full message queue, while the second one is for serialising the *message writes*. The size of the message is found when the message queue is generated.

Out of all these, shared memory is supposed to be the fastest IPC mechanism. As we discussed, this IPC allows processes to share the physical memory. MMS (memory management system) does the creation of the shared physical memory regions. The system call `sys_shmat()` links the shared pages to the user processes virtual memory space.

Data structures (for IPC mechanisms)

Now, we can talk about the data structures to implement these IPC mechanisms. The implementation of signals is done using the signal field in the *task_struct* structure. And every signal is conveyed using a bit in the field. So the total number of signals is limited to the number of bits in a word.

The *wait_queue* structure, which has a pointer to the associated *task_struct*, is the one linked to wait queues. *Inode* is the filesystem for pipes. It records the extra pipe-specific information in the *pipe_inode_info*. And it has a wait queue (for the processes), the quantity of data, and the number of

processes that are reading (/writing) from (/to) the pipe.

For *file locks* we have *file_lock* structure. This has a pointer to a *task_struct* for the owning process, a file descriptor of the locked file, and a wait queue for processes and the area where it is locked. For every open file, the *file_lock* structures will link to a list. And socket data structure is used for representing the UNIX domain sockets.

When it comes to the system V IPC objects, we can see that these are created in the kernel. The associated access permissions are mentioned in the *ipc_perm* structure. Semaphores are represented with the *sem* structure (which also has the value and the PID of the process that used semaphore). The arrays are done using the *semid_ds* structure (that carries access permissions and the time of the last operation, the pointer to the first semaphore in the list, etc). The *Sem_undo* (structure) is employed to get an array of semaphore operations performed by a process (used during the killing of the process).

Message queues are represented by the *msqid_ds* structure, which also has the control and management information. It can store the following:

- access permissions
- the current number of bytes in the queue
- the number of messages
- the size of the queue
- the process number of the last message sent
- the process number of the last message received
- the time of the last change and message sent and received
- link fields to implement the message queue

And a message is recorded in the kernel with a *msg* structure. This has information about the link field (for the link list), the message type, address of data (message data) and the length. The *Shmid_ds* structure is used to represent the shared memory implementation. It has access control permissions, PIDs of the creator, number of processes to which the shared physical memory region (and number of pages that make this zone) is linked, detach and change times, etc.

We can now summarise the overall architecture using Figure 4. (You can see that this is a very generalised view to just put the idea in a nutshell; for the exact structure, please refer to the earlier columns.)

Newbie zone

I'm going to discuss a few things that I have been telling novice programmers. First of all, let's look at the differences between Linux and the UNIX kernel. Since Linux is based on the UNIX architecture, there are many similarities. But Linux is not just a copy of UNIX.

The most interesting feature of Linux is that it allows you to load modules dynamically. We have already learned how to write a module and load that to the kernel. This feature is remarkable especially when you find that the kernel is monolithic in nature, but can still support this one feature. Another characteristic that needs special emphasis is the treatment of threads in Linux. It views threads and normal processes with the same eye. It has symmetrical

multi-processor (SMP) support (though this is available in most of the modern UNIX derivatives, but it is not present in the traditional ones) and is pre-emptive. Thus you can pre-empt any task even if it is running in the kernel. You might have seen this in derivatives like IRIX and Solaris, but not in traditional ones.

The operating system supports an object-oriented device model and sysfs (user-space device filesystem). And you can also see that features like STREAMS are not available in Linux. Linux designers avoided many such poorly deployed ideas. What else? It is free and open source!

Tips

1. You can join the Linux kernel mailing list. To subscribe, go to <http://vger.kernel.org>. One thing that you need to take care of is that when you subscribe to lists like `linux-kernel`, `vger.kernel.org`, please create a new folder in your e-mail client (or create a filter, if you are in Gmail). Since this is a very high traffic list, you will get about 200 messages per day! By subscribing, you can keep yourself updated and interact with other developers.
2. I have already mentioned the ways to download the kernel source to your system (both by a direct download and command-line based). You can find the source under `/usr/src/your_linux_kernel_version`. You don't need to use this when you are performing some edits. You can maintain a directory under your *home* folder and do the trials. You need to switch to root (`su -`) only when you are about to install it. Also, when you move from one version to another you don't have to go for a full upgrade, but use patches. For an incremental patch, you can run the code below (after getting inside your source tree):

```
$ patch -p1 < ../patch-a.b.c
```

Normally, we apply this against the preceding version of the kernel that you want to use.

3. There are many utilities that support you while programming and debugging. There are even GUI versions available for many programs. The *make config* (command line) utility allows you to do the modifications (customisation) to the kernel (please refer to Day 14). You can also try a *ncurses-based* graphical utility like *menuconfig* for this. *Make xconfig* (or try *gconfig* if you need a gtk+ based one) is another X11-based graphical utility that aids you. You can make use of all these for customisation options.

By using *make defconfig* you can make a config based on the default setting that suits your architecture. This is a very good tool for beginners since you may not have configured the kernel earlier:

```
root@GNU-BOX:/usr/src/linux-source-2.6.27# make defconfig
HOSTCC scripts/basic/fixdep
HOSTCC scripts/basic/docproc
```

```
HOSTCC scripts/kconfig/conf.o
scripts/kconfig/conf.c: In function 'conf_askvalue':
scripts/kconfig/conf.c:104: warning: ignoring return value of 'fgets', declared
with attribute warn_unused_result
scripts/kconfig/conf.c: In function 'conf_choice':
scripts/kconfig/conf.c:306: warning: ignoring return value of 'fgets', declared
with attribute warn_unused_result
HOSTCC scripts/kconfig/kxgettext.o
SHIPPED scripts/kconfig/zconf.tab.c
SHIPPED scripts/kconfig/ex.zconf.c
SHIPPED scripts/kconfig/zconf.hash.c
HOSTCC scripts/kconfig/zconf.tab.o
In file included from scripts/kconfig/zconf.tab.c:2486:
scripts/kconfig/confdata.c: In function 'conf_write':
scripts/kconfig/confdata.c:501: warning: ignoring return value of 'fwrite',
declared with attribute warn_unused_result
scripts/kconfig/confdata.c: In function 'conf_write_autoconf':
scripts/kconfig/confdata.c:739: warning: ignoring return value of 'fwrite',
declared with attribute warn_unused_result
scripts/kconfig/confdata.c:740: warning: ignoring return value of 'fwrite',
declared with attribute warn_unused_result
In file included from scripts/kconfig/zconf.tab.c:2487:
scripts/kconfig/expr.c: In function 'expr_print_file_helper':
scripts/kconfig/expr.c:1090: warning: ignoring return value of 'fwrite',
declared with attribute warn_unused_result
HOSTLD scripts/kconfig/conf
*** Default configuration is based on 'i386_defconfig'
#
# configuration written to .config
#
```

After executing this, you can find the *.config* file in your source tree with entries like this:

```
# Automatically generated make config: don't edit
# Linux kernel version: 2.6.27.18
# Tue Aug 12 08:39:24 2009
#
# CONFIG_64BIT is not set
CONFIG_X86_32=y
# CONFIG_X86_64 is not set
CONFIG_X86=y
CONFIG_ARCH_DEFCONFIG="arch/x86/configs/i386_defconfig"
```

make oldconfig is another utility that helps you in validating and updating the configuration.

Please note that if you are on an earlier kernel, *make dep* may be required. You can see that I am now trying a more recent version. There are a few hurdles when you change the versions. I remember one reader had an issue when he used macros instead of syscall. For version 2.6.18 onwards, the API has been changed and you need to use syscall (*'...performs the system call whose assembly language interface has the specified number with the specified arguments.'*). More details are available at: www.kernel.org/doc/man-pages/online/pages/man2/

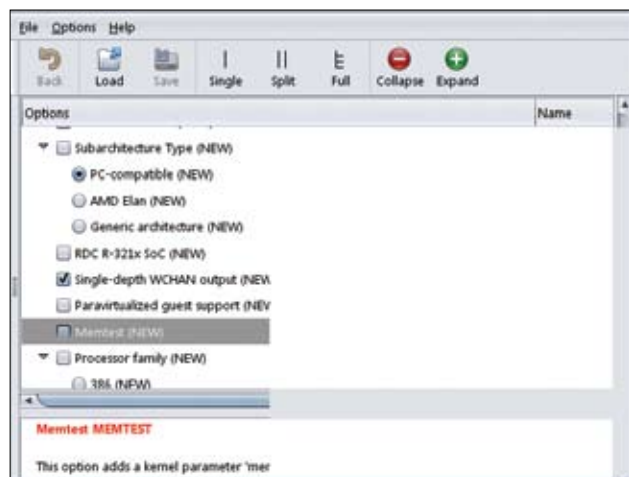


Figure 5: Kernel configuration using the graphical gconf tool

`syscall.2.html`

4. To minimise the build noise, you can use the following command:

```
make > ../some_out
```

This will just direct the output from *make* (but you can still see the errors). I would suggest that you redirect them to `/dev/null`.

5. *An advanced feature:* The *make* program allows you to split the entire build process into a number of *jobs*. So, if your processor is quite efficient, you can split the process and run the jobs concurrently. This is a recommended way for hackers, as this will reduce your I/O wait time. You can just use the command given below:

```
make -jn
```

...where *n* is the number of jobs. You may choose this (value of *n*) based on your processor.

When it comes to the kernel installation side, it is better that you follow your architecture and boot loader specific instructions. But for modules you can simply issue the command below as the root:

```
make modules_install
```

6. Novices can use the files like the *package-list*, which may be located in the root of the source (for example, `/linux-source-kernel_version/package-list`) to comprehend the files included, dependency (kernel image and modules), etc.

Package: acpi-modules
Depends: kernel-image
Priority: standard
Description: Support for ACPI

Package: fat-modules
Depends: kernel-image

Priority: standard

Description: FAT filesystem support

This includes Windows FAT and VFAT support.

Package: fb-modules

Depends: kernel-image

Priority: standard

Description: Framebuffer modules

Back to action

When you are programming, you can find that the kernel is not linked to the standard C library, mostly due to the size and speed factors (as even a *decent subset* of the C library is quite large for the kernel). But you can use many *libc* functions (say, functions for string manipulation) as they are implemented inside the kernel.

Now, you might ask about the *printf()* in C. Of course, the kernel can't access *printf()* but it has access to *printk()*. It uses the *syslog* program to read the log buffer entries made by *printk()* and its (function's) usage is similar to that of *printf()*.

Process descriptor and the task structure

Task list (a circular doubly-linked list) is used by the kernel to record the list of processes.

Each element is a process descriptor of type struct *task_struct* and the reference can be found in the `/linux/sched.h`:

```
extern void proc_sched_show_task(struct task_struct *p, struct seq_file *m);
extern void proc_sched_set_task(struct task_struct *p);
extern void
print_cfs_rq(struct seq_file *m, int cpu, struct cfs_rq *cfs_rq);
#else
static inline void
proc_sched_show_task(struct task_struct *p, struct seq_file *m)
{
}
static inline void proc_sched_set_task(struct task_struct *p)
{
}
static inline void
print_cfs_rq(struct seq_file *m, int cpu, struct cfs_rq *cfs_rq)
{
}
#endif

extern unsigned long long time_sync_thresh;
```

The process descriptor has all the data concerning a specific process. The *task_struct* structure is allocated using the slab allocator (to enable cache colouring and object reuse). If you look at Figure 6, you can see the struct *thread_info*. It remains on the top and bottom of it, for the stacks that grow up. The following code shows the *thread_info* structure defined on an *x86* as:

```
struct thread_info {
```

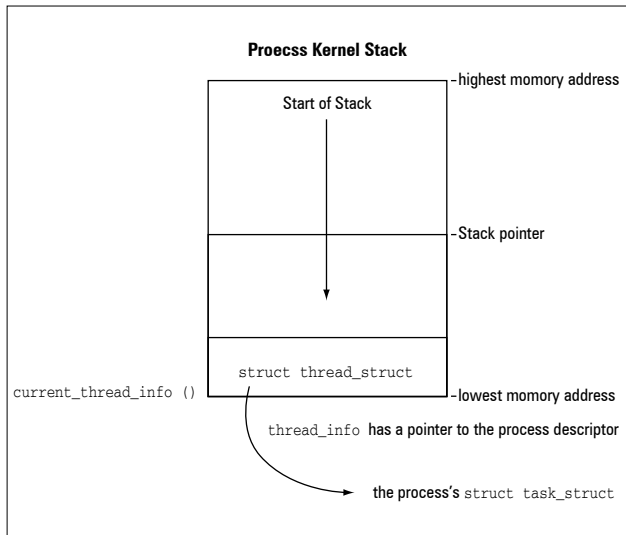


Figure 6: Process descriptor and stack

```

struct task_struct *task;
struct exec_domain *exec_domain;
unsigned long flags;
unsigned long status;
__u32 cpu;
__s32 preempt_count;
mm_segment_t addr_limit;
struct restart_block restart_block;
unsigned long previous_esp;
__u8 supervisor_stack[0];
};

```

You may note that the task element of the structure is a pointer to a task's *task_struct*.

Runqueues

The fundamental data structure in the scheduler is *runqueue*. For its definition, please see the code (commented) below:

```

struct runqueue {
    spinlock_t lock; /* spin lock that protects this runqueue */
    unsigned long nr_running; /* number of runnable tasks */
    unsigned long nr_switches; /* context switch count */
    unsigned long expired_timestamp; /* time of last array swap */
    unsigned long nr_uninterruptible; /* uninterruptible tasks */
    unsigned long long timestamp_last_tick; /* last scheduler tick */
    struct task_struct *curr; /* currently running task */
    struct task_struct *idle; /* this processor's idle task */
    struct mm_struct *prev_mm; /* mm_struct of last ran task */
    struct prio_array *active; /* active priority array */
    struct prio_array *expired; /* the expired priority array */
    struct prio_array arrays[2]; /* the actual priority arrays */
    struct task_struct *migration_thread; /* migration thread */
    struct list_head migration_queue; /* migration queue */
    atomic_t nr_iowait; /* number of tasks waiting on I/O */
};

```

If you wish to see the complete code, please look at the *kernel/sched.c* (see struct *runqueue* in the code) file. The *real-timeclasses* related field in a *runqueue* is given here for your reference:

```

struct rt_rq {
    struct rt_prio_array active;
    unsigned long rt_nr_running;
#if defined CONFIG_SMP || defined CONFIG_RT_GROUP_SCHED
    int highest_prio; /* highest queued rt task prio */
#endif
#ifdef CONFIG_SMP
    unsigned long rt_nr_migratory;
    int overloaded;
#endif
    int rt_throttled;
    u64 rt_time;
    u64 rt_runtime;
    /* Nests inside the rq lock: */
    spinlock_t rt_runtime_lock;

#ifdef CONFIG_RT_GROUP_SCHED
    unsigned long rt_nr_boosted;

    struct rq *rq;
    struct list_head leaf_rt_rq_list;
    struct task_group *tg;
    struct sched_rt_entity *rt_se;
#endif
};


```

```

#ifdef CONFIG_SMP

```

Runqueue is basically a list of processes that are scheduled to run in a processor, and there will be one such list for each processor. But a given process will not appear in two lists. You might have noticed that I have referred to a .c file and not a header one. If you call a header file, it can include codes outside of the scheduler. In order to prevent this from happening, we refer to the other file. Two macros are important when it comes to *runqueue*, and they are the macro *cpu_rq* (which returns the pointer to the *runqueue* linked to a particular processor) and the *this_rq()* macro (which points to the current one). There are other macros also that are linked to this (like *task_rq(task)*).

We are done for today. Wait for the next instalment to continue the voyage. Happy kernel hacking! 

By: Aasis Vinayak PG

The author is a hacker and a free software activist who does programming in the open source domain. He is the developer of V-language—a programming language that employs AI and ANN. His research work/publications are available at www.aasisvinayak.com

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

HIGHLIGHTS

- A cost-effective marketing tool
- A user-friendly format for customers to contact you
- A dedicated section with yellow back-ground, and hence will stand out
- Reaches to tech-savvy IT implementers and software developers
- 80% of LFY readers are either decision influencers or decision takers
- Discounts for listing under multiple categories
- Discounts for booking multiple issues

FEATURES

- Listing is categorised on the basis of products and services
- Complete contact details plus 30-word description of organisation
- Option to print the LOGO of the organisation too (extra cost)
- Option to change the organisation description for listings under different categories



TARIFF

Category Listing

ONE Category	Rs 2,000
TWO Categories.....	Rs 3,500
THREE Categories.....	Rs 4,750
ADDITIONAL Category	Rs 1,000

Value-add Options

LOGO-plus-Entry.....	Rs 500
Highlight Entry (white background).....	Rs 1,000
Per EXTRA word (beyond 30 words).....	Rs 50

KEY POINTS

- Above rates are per-category basis.
- Above rates are charges for publishing in a single issue of LFY.
- Max. No. of Words for Organisation Description: 30 words.

TERMS & CONDITIONS

- Fill the form (below).
- You can use multiple copies of the form for multiple listings under different categories.
- Payment to be received along with booking.

✂ Tear & Send

ORDER FORM

Tear & Send ✂

Organisation Name (70 characters): _____

Description (30 words): _____

Email: _____ Website: _____

STD Code: _____ Phone: _____ Mobile: _____

Address (will not be published): _____

City/Town: _____ Pin-code: _____

CATEGORIES

- | | | |
|---|--|---|
| <input type="checkbox"/> CONSULTANTS | <input type="checkbox"/> HIGH PERFORMANCE COMPUTING | <input type="checkbox"/> SOFTWARE DEVELOPMENT |
| <input type="checkbox"/> CONSULTANT (FIRM) | <input type="checkbox"/> IT INFRASTRUCTURE SOLUTIONS | <input type="checkbox"/> TRAINING FOR PROFESSIONALS |
| <input type="checkbox"/> EMBEDDED SOLUTIONS | <input type="checkbox"/> LINUX-BASED WEB-HOSTING | <input type="checkbox"/> TRAINING FOR CORPORATE |
| <input type="checkbox"/> ENTERPRISE COMMUNICATION SOLUTIONS | <input type="checkbox"/> MOBILE SOLUTIONS | <input type="checkbox"/> THIN CLIENT SOLUTIONS |

Please find enclosed a sum of Rs. _____ by DD/ MO//crossed cheque* bearing the No. _____ dt. _____ in favour of

EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs. 50 on non-metro cheque) towards the cost of _____ FOSS Yellow Pages advertisement(s)

or charge my credit card ☐ VISA ☐ Master Card Please charge Rs. _____

against my credit card No. _____

C V V No. _____ (Mandatory)

Date of Birth ____ / ____ / ____ (dd/mm/yy) Card Expiry Date ____ / ____ (mm/yy)



EFY Enterprises Pvt Ltd., D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020
Ph: 011-26810601-03, Fax: 011-26817565, Email: info@efyindia.com; Website: www.efyindia.com

Signature (as on the card)

To Book Your Listing, Call: Dhiraj (Delhi: 0981206582), Somaiah (B'lore: 09986075717)

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

Consultant (Firm)

IB Services

Free Installation of GNU/Linux on Laptops and Desktops. Thin client solutions based on Debian and Ubuntu. Laptops and Desktops pre-installed with Debian and Ubuntu. Migration to GNU/Linux. Data Recovery. Navi Mumbai
Kerala
Mobile: 09847446918
Email: ibmanoj@gmail.com
Web : www.ibservices.in

OS3 Infotech

- Silver Solutions Partner for Novell
- High Availability Computing Solutions
- End-to-end Open Source Solutions Provider
- Certified Red Hat Training Partner
- Corporate and Institutional Training

Navi Mumbai
Mobile: 09324113579
Email: info@os3infotech.com
Web: www.os3infotech.com

Taashee Linux Services

100% Support on LINUX, OSS & JBOSS related projects. We specialize in high-availability and high-performance clusters, remote and onsite system management, maintenance services, systems planning, Linux & JBOSS consulting & Support services.
Hyderabad
Mobile: 09392493753, Fax: 040-40131726
Email: manojkummar@taashee.com
Web: www.taashee.com

Torrid Networks Pvt Ltd

Torrid is a leading provider of information technology consulting with focus on Information Security Services and Open Source Solutions.
Noida
Tel: 0120-4545100, Mobile: 09015505583
Fax: 0120-4235064
Email: tarun.kassana@torridnetworks.com
Web: www.torridnetworks.com

Computer (UMPC) For Linux And Windows

Comptek International

World's smallest computer comptek wibrain B1 UMPC with Linux, Touch Screen, 1 GB RAM 60GB, Wi-Fi, Webcam, upto 6 hour battery (opt.), USB Port, max 1600x1200 resolution, screen 4.8", 7.5"x3.25" Size, weight 526 gm.
New Delhi

Mobile: 09968756177, Fax: 011-26187551
Email: comptekdelhi@compteki.com
Web: www.compteki.com
or www.compteki.in

Education & Training

Aptech Limited

IT, Multimedia and Animation Education and Training
Mumbai
Tel: 022-28272300, 66462300
Fax: 022-28272399
Email: customercare@aptech.ac.in
Web: www.aptech-education.com, www.arena-multimedia.com

IT-Campus: Academy of Information Technology

IT training and solution company with over 12 years of experience. - RHCE
• Software Training • Hardware Training
• Multimedia And Animation • Web Designing • Financial Accounting
Kota (Raj.)
Tel: 0744-2503155, Mobile: 09828503155
Fax: 0744-2505105
Email: m_trilok@yahoo.com
Web: www.doeacc4u.com

Mahan Computer Services (I) Limited

Established in 1990, the organization is primarily engaged in Education and Training through its own & Franchise centres in the areas of IT Software, Hardware, Networking, Retail Management and English. The institute also provides customized training for corporates.
New Delhi
Tel: 011-25916832-33
Email: info@mahanindia.com
Web: www.mahanindia.com

Enterprise Comm. Solutions

Aware Consultants

We specialize in building and managing Ubuntu/Debian Linux servers and provide good dependable system administration. We install and maintain in-house corporate servers. We also provide dedicated and shared hosting as well as reliable wireless/hybrid networking.
Bangalore
Tel: 080-26724324
Email: sales@aware.co.in
Web: www.aware.co.in

Cynapse India Private Limited

We are the creators of open source

product cyn.in. cyn.in is a web 2.0 group collaboration software created by Cynapse, that inter-connects your people with each other and their collective knowledge, seamlessly. It combines the capabilities of collaboration tools like wikis, blogs, file repositories, micro blogs, discussions, audio, videos, and other social applications into a seamless platform. cyn.in helps teams to build collaborative knowledge by sharing and discussing various forms of digital content within a secure, unified application that is accessible using a web based interface or a rich desktop client.
Mumbai

Tel: 022-28445858, 28445629
Email: lisa@cynapse.com
Web: www.cynapse.com

DeepRoot Linux Pvt Ltd

Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom We believe in: sharing, compassion and plain action. Backed by full-time hackers. Quick deployment, easy management. Guaranteed.
Bangalore
Tel: 080-40890000
Email: start@deeproot.co.in
Web: www.deeproot.in

ESQUBE Communications Solutions Pvt Ltd

Founders of ESQUBE are faculty at the Indian Institute of Science, Bangalore and carry over eight decades of experience and fundamental knowledge in the field of DSP and Telecommunication. ESQUBE plays a dominant role in the creation of IP in the domain of Sensors, Signals and Systems.
Bangalore
Tel: 080-23517063
Email: info@esqube.com
Web: www.esqube.com

Keen & Able Computers Pvt Ltd

Microsoft Outlook compatible open source Enterprise Groupware Mobile push, Email Syncing of Contacts/Calendar/Tasks with mobiles
• Mail Archival • Mail Auditing • Instant Messaging
New Delhi
Tel: 011-30880046, 30880047
Mobile: 09810477448, 09891074905
Email: info@keenable.com
Web: www.keenable.com

netCORE
The Innovation Company

Netcore Solutions Pvt Ltd

No.1 company for providing Linux Based Enterprise Mailing solution with around 1500+ Customer all over India. Key Solutions:
• Enterprise Mailing and Collaboration Solution • Hosted Email Security • Mail Archiving Solution • Push Mail on Mobile • Clustering Solution
Mumbai
Tel: 022-66628000
Mobile: 09322985222
Email: kalpit@netcore.co.in
Web: www.netcore.co.in

 **redhat.**

Red Hat India Pvt Ltd

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.
Mumbai
Tel: 022-39878888
Email: marketing-in@redhat.com
Web: www.redhat.in

Hardware & Networking Institute

Xenitis Technolab Pvt Ltd

Xenitis Technolab is the first of its kind, state-of-the-art infrastructure, Hardware, Networking and I.T Security training institution headquartered in Kolkata. Technolab is the training division of Xenitis group of Companies. It is the proud owner of 'Aamar PC', the most popular Desktop brand of Eastern India. These ranges of PC's are sold in the west under the brand name of 'Aamchi PC', in the north as 'Aapna PC' and in the south as 'Namma PC'.
Kolkata
Tel: 033-22893280
Email: srinku@xenitisgroup.com
Web: www.technolabindia.com

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

IT Infrastructure Solutions

Absolut Info Systems Pvt Ltd

Open Source Solutions Provider.
Red Hat Ready Business Partner.
Mail Servers/Anti-spam/GUI
interface/Encryption, Clustering &
Load Balancing - SAP/Oracle/Web/
Thin Clients, Network and Host
Monitoring, Security Consulting,
Solutions, Staffing and Support.

New Delhi

Tel: +91-11-26494549

Fax: +91-11-4175 1823

Mobile: +91-9873839960

Email: sales@aisplglobal.com

Web: www.aisplglobal.com

Advent Infotech Pvt Ltd

Advent has an experienced techno-
marketing team with several years of
experience in Networking & Telecom
business, and is already making
difference in market place. ADVENT
qualifies more as Value Added
Networking Solution Company, we
offers much to customers than just
Routers, Switches, VOIP, Network
Management Software, Wireless
Solutions, Media Conversion, etc.

New Delhi

Tel: 46760000, 09311166412

Fax: 011-46760050

Email: marketingsupport@

adventelectronics.com

Web: www.adventelectronics.com

Asset Infotech Ltd

We are an IT solution and training
company with an experience of 14
years, we are ISO 9001: 2000. We
are partners for Red-Hat, Microsoft,
Oracle and all Major software
companies. We expertise in legal
software ans solutions.

Dehradun

Tel: 0135-2715965, Mobile: 09412052104

Email: piyush@asset.net.in

Web: www.asset.net.in

BakBone Software Inc.

BakBone Software Inc. delivers
complexity-reducing data protection
technologies, including award-
winning Linux solutions; proven
Solaris products; and application-
focused Windows offerings that
reliably protect MS SQL, Oracle,
Exchange, MySQL and other
business critical applications.

New Delhi

Tel: 011-42235156

Email: ashish.gupta@bakbone.com

Web: www.bakbone.com

Clover Infotech Private Limited

Clover Infotech is a leading
technology services and solutions
provider. Our expertise lies in
supporting technology products
related to Application, Database,
Middleware and Infrastructure. We
enable our clients to optimize their
business through a combination of
best industry practices, standard
processes and customized client
engagement models. Our core
services include Technology
Consulting, Managed Services and
Application Development Services.

Mumbai

Tel: 022-2287 0659, Fax: 022-2288 1318

Mobile: +91 99306 48405

Email: business@cloverinfotech.com

Web: www.cloverinfotech.com

DeepRoot Linux Pvt Ltd

Pure & Exclusive Free Software
Business. Creators of the deepOfix
Mail Server. We provide: airtight
solutions, solid support and Freedom
We believe in: sharing, compassion
and plain action. Backed by full-time
hackers. Quick deployment, easy
management. Guaranteed.

Bangalore

Tel: 080-40890000

Email: start@deeproot.co.in

Web: www.deeproot.in

Duckback Information Systems Pvt Ltd

A software house in Eastern India.
Business partner of Microsoft,
Oracle, IBM, Citrix , Adobe, Redhat,
Novell, Symantec, McAfee, Computer
Associates, Veritas , Sonic Wall

Kolkata

Tel: 033-22835069, 9830048632

Fax: 033-22906152

Email: asis@duckback.com

Web: www.duckback.co.in

HBS System Pvt Ltd

System Integrators & Service
Provider.Partner of IBM, DELL, HP,
Sun, Microsoft, Redhat, Trend Micro,
Symentic Partners of SUN for their
new startup E-commerce initiative
Solution Provider on REDHAT,
SOLARIS & JAVA

New Delhi

Tel: 011-25767117, 25826801/02/03

Fax: 25861428

Email: amitall@hbsindia.com

Ingres Corporation

Ingres Corporation is a leading
provider of open source database
software and support services.
Ingres powers customer success
by reducing costs through highly
innovative products that are hallmarks
of an open source deployment and
uniquely designed for business critical
applications. Ingres supports its
customers with a vibrant community
and world class support, globally.
Based in Redwood City, California,
Ingres has major development, sales,
and support centers throughout
the world, and more than 10,000
customers in the United States and
internationally.

New Delhi

Tel: 011-40514199, Fax: +91 22 66459537

Email: sales@ingres.com; info@ingres.com

Web: www.ingres.com

Keen & Able Computers Pvt Ltd

Open Source Solutions Provider.
Red Hat Ready Business Partner.
Mail Servers/Anti-spam/GUI
interface/Encryption, Clustering &
Load Balancing - SAP/Oracle/Web/
Thin Clients, Network and Host
Monitoring, Security Consulting,
Solutions, Staffing and Support.

New Delhi-110019

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074905

Email: info@keenable.com

Web: www.keenable.com

LDS Infotech Pvt Ltd

Is the authorised partner for RedHat
Linux, Microsoft, Adobe, Symantec,
Oracle, IBM, Corel etc. Software
Services Offered: •Collaborative
Solutions •Network Architecture
•Security Solutions •Disaster
Recovery •Software Licensing
•Antivirus Solutions.

Mumbai

Tel: 022-26849192

Email: sales@ldsinfotech.com

Web: www.ldsinfotech.com

Pacer Automation Pvt Ltd

Pacer is leading providers of IT
Infrastructure Solutions. We are
partners of HP, Redhat, Cisco,
Vmware, Microsoft and Symantec.
Our core expertise exists in,
Consulting, building and Maintaining
the Complete IT Infrastructure.

Bangalore

Tel: 080-42823000, Fax: 080-42823003

Email: solutions@pacerautomation.com

Web: www.pacerautomation.com



Red Hat India Pvt Ltd

Red Hat is the world's leading
open source solutions provider.
Red Hat provides high-quality,
affordable technology with its
operating system platform, Red
Hat Enterprise Linux, together with
applications, management and
Services Oriented Architecture (SOA)
solutions, including JBoss Enterprise
Middleware. Red Hat also offers
support, training and consulting
services to its customers worldwide.

Mumbai

Tel: 022-39878888

Email: marketing-in@redhat.com

Web: www.redhat.in

Srijan Technologies Pvt Ltd

Srijan is an IT consulting company
engaged in designing and
building web applications, and IT
infrastructure systems using open
source software.

New Delhi

Tel: 011-26225926, Fax: 011-41608543

Email: business@srijan.in

Web: www.srijan.in

TechnInfotech™

A company focussed on Enterprise
Solution using opensource software.
Key Solutions:

- Enterprise Email Solution
- Internet Security and Access
Control
- Managed Services for Email
Infrastructure.

Mumbai

Tel: 022-66338900; Extn. 324

Email: sales@techninfotech.com

Web: www.techninfotech.com

To advertise in this section,
please contact

Somaiah (B'lore):

09986075717) Dhiraj (Delhi):

09811206582) on

011-2681-0602 Extn. 222

FOSS Yellow Pages

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

Tetra Information Services Pvt Ltd

One of the leading open source providers. Our cost effective business ready solutions caters of all kind of industry vehicles.

New Delhi

Tel: 011-46571313, Fax: 011-41620171

Email: sales@tetra.in

Web: www.tetra.in

Tux Technologies

Tux Technologies provides consulting and solutions based on Linux and Open Source software. Focus areas include migration, mail servers, virus and spam filtering, clustering, firewalls, proxy servers, VPNs, server optimization.

New Delhi

Tel: 011-27348104, Mobile: 09212098104

Email: info@tuxtechnologies.co.in

Web: www.tuxtechnologies.co.in

Veeras Infotek Private Limited

An organization providing solutions in the domains of Infrastructure Integration, Information Integrity, Business Applications and Professional Services.

Chennai

Tel: 044-42210000, Fax: 28144986

Email: info@veeras.com

Web: www.veeras.com

Linux-Based Web-Hosting

Manas Hosting

Manas-Hosting is a Bangalore-based company that is dedicated in helping small and midsize business companies to reach customers online. We believe that by creating a website, all you have is just web presence; but to get effective traffic on your website, it is equally important to have a well designed one. This is why we provide the best of Web Hosting and Web Designing services. Also, our services are backed with exceptionally good quality and low costs

Bangalore

Tel: 080-42400300

Email: enquiry@manashosting.com

Web: www.manashosting.com

Linux Desktop

Indserve Infotech Pvt Ltd

OpenLx Linux with Calcutate (Financial Accounting & Inventory on Linux) offers a complete Linux

Desktop for SME users. Its affordable (Rs. 500 + tax as special scheme), Friendly (Graphical UserInterface) and Secure (Virus free).

New Delhi

Tel: 011-26014670-71, Fax: 26014672

Email: info@openlx.com

Web: www.openlx.com

Linux Experts

Intaglio Solutions

We are the training and testing partners of Red-Hat and the first to conduct RHCCS exam in delhi for the first time ever.

New Delhi

Tel: 011-41582917, 45515795

Email: info@intaglio-solutions.com

Web: www.intaglio-solutions.com

Linux Vendor/Distributors

GT Enterprises

Authorized distributors for Red Hat and JBoss range of products. We also represent various OS's Applications and Developer Tools like SUSE, VMWare, Nokia Qt, MySQL, Codeweavers, Ingres, Sybase, Zimbra, Zend-A PHP Company, High Performance Computing Solutions from The Portland Group, Absoft, Pathscale/Qlogic and Intel Compilers, Scalix-Messaging solution on Linux Platform.

Bangalore

Mobile: +91-9845009939, +91-9343861758

Email : sales@gte-india.com

Web: www.gte-india.com

Taurusoft

Contact us for any Linux Distribution at reasonable rates. Members get additional discounts and Free CD/DVDs with each purchase. Visit our website for product and membership details

Mumbai

Mobile: 09869459928, 09892697824

Email: taurusoft@gmail.com

Web: www.taurusoft.netfirms.com

Software Subscriptions

Blue Chip Computers

Available Red Hat Enterprise Linux, Suse Linux Enterprise Server / Desktop, JBoss, Oracle, ARCserve Backup, AntiVirus for Linux, Verisign/Thawte/GeoTrust SSL Certificates

and many other original software licenses.

Mumbai

Tel: 022-25001812, Mobile: 09821097238

Email: bluechip@vsnl.com

Web: www.bluechip-india.com

Software Development

Carizen Software (P) Ltd

Carizen's flagship product is Rainmail Intranet Server, a complete integrated software product consisting modules like mail sever, proxy server, gateway anti-virus scanner, anti-spam, groupware, bandwidth aggregator & manager, firewall, chat server and fax server. Infrastructure.

Chennai

Tel: 044-24958222, 8228, 9296

Email: info@carizen.com

Web: www.carizen.com

DeepRoot Linux Pvt Ltd

Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom We believe in: sharing, compassion and plain action. Backed by full-time hackers. Quick deployment, easy management. Guaranteed.

Bangalore

Tel: 080-40890000

Email: start@deeproot.co.in

Web: www.deeproot.in



InfoAxon Technologies Ltd

InfoAxon designs, develops and supports enterprise solutions stacks leveraging open standards and open source technologies. InfoAxon's focus areas are Business Intelligence, CRM, Content & Knowledge Management and e-Learning.

Noida

Tel: 0120-4350040, Mobile: 09810425760

Email: sales@infoaxon.com

Web: http://opensource.infoaxon.com

Integra Micro Software Services (P) Ltd

Integra focuses on providing professional services for software development and IP generation to customers. Integra has a major practice in offering Telecom Services and works for Telecom companies, Device Manufacturers, Networking

companies, Semiconductor and Application development companies across the globe.

Bangalore

Tel: 080-28565801/05, Fax: 080-28565800

Email: tpvarun@integramicro.com

Web: www.integramicroservices.com

iwebtune.com Pvt Ltd

iwebtune.com is your one-stop, total web site support organisation. We provide high-quality website services and web based software support to any kind of websites, irrespective of the domain or the industry segments.

Bangalore

Tel: 080-4115 2929

Email: santosh@iwebtune.com

Web: www.iwebtune.com

Sarvasv Technologies Pvt Ltd

Web OSX Windows Linux Mobile LAMP W3C apps. Agile, BDD, TDD development. On-Demand, SaaS, ERP, CRM, HRM. Open Source integration. Online project management. VCS (git, svn). automated deployment.

Greater Noida

Mobile: 09711472457, 09810270848

Email: sales@sarvasv.in

Website: www.sarvasv.in

Unistal Systems Pvt Ltd

Unistal is pioneer in Data Recovery Software & Services. Also Unistal is national sales & support partner for BitDefender Antivirus products.

New Delhi

Tel: 011-26288583, Fax: 011-26219396

Email: isales@unistal.com

Web: www.unistal.com

Software and Web Development

Bean eArchitect Integrated Services Pvt Ltd

Application Development, Web Design, SEO, Web Marketing, Web Development.

Navi Mumbai

Tel: 022-27821617, Mobile: 9820156561

Fax: 022-27821617

Email: infodesk@beanarchitect.com

Web: www.beanarchitect.com

Mr Site Takeaway Website Pvt Ltd

Our product is a unique concept in India using which a person without having any technical knowledge can create his website within 1 hour; we also have a Customer Care Center

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

In India for any kind of after sales help. We are already selling it world over with over 65,000 copies sold. It comes with FREE Domain Name, Web Hosting and Customer Care Center for Free Support via Phone and Email and features like PayPal Shopping Cart, Guestbook, Photo Gallery, Contact Form, Forums, Blogs and many more. The price of complete package is just Rs 2,999 per year.

Patiala

Mobile: 91-9780531682

Email: pardeep@mrsite.co.in

Web: www.mrsite.co.in

Salah Software

We are specialized in developing custom strategic software solutions using our solid foundation on focused industry domains and technologies. Also providing superior Solution Edge to our Clients to enable them to gain a competitive edge and maximize their Return on Investments (ROI).

New Delhi

Tel: 011-41648668, 66091565

Email: ceo@salahsoftware.com

Web: www.salahsoftware.com

Thin Client Solutions

Digital Waves

The 'System Integration' business unit offers end-to-end Solutions on Desktops, Servers, Workstations, HPC Clusters, Render Farms, Networking, Security/Surveillance & Enterprise Storage. With our own POWER-X branded range of Products, we offer complete Solutions for Animation, HPC Clusters, Storage & Thin-Client Computing

Mobile: 09880715253

Email: ranga@digitalwaves.in

Web: www.digitalwaves.in

Enjay Network Solutions

Gujarat based ThinClient Solution Provider. Providing Small Size ThinClient PCs & a Full Featured ThinClient OS to perfectly suite needs of different working environment. Active Dealer Channel all over India.

Gujarat

Tel: 0260-3203400, 3241732, 3251732,

Mobile: 09377107650, 09898007650

Email: info@enjayworld.com

Web: www.enjayworld.com

Training for Corporate

Bascom Bridge

Bascom Bridge is Red Hat Certified partner for Enterprise Linux 5 and also providing training to the individuals and corporate on other open source technologies like PHP, MySQL etc.

Ahmedabad

Tel: 079-27545455—66

Fax: 079-27545488

Email: info@bascombridge.com

Web: www.bascombridge.com

Brainnet

Kolkata

Tel: 033-40076450

Email: brainnet@brainware-india.com

Web: www.brainware-india.com

Centre for Excellence in Telecom Technology and Management (CETTM), MTNL

MTNL's Centre for Excellence in Telecom Technology and Management (CETTM) is a state of the art facility to impart Technical, Managerial and corporate training to Telecom; Management personnel. CETTM has AC lecture halls, computer Labs and residential facility.

Mumbai

Tel: 022-25714500, 25714586, 25714585, 25714586

Fax: 022-25706700

Email: contact@cetm.mtnl.in

Web: http://cetm.mtnl.in/infra

Complete Open Source Solutions

RHCT, RHCE and RHCSA training.

Hyderabad

Tel: 040-66773365, 9849742065

Email: nayak.sujeet@gmail.com

Web: www.cossindia.com

ElectroMech

Redhat Linux and open source solution, RHCE, RHCSA training and exam center, Ahmedabad and Vadodara

Ahmedabad

Tel: 079-40027898

Email: electromech@electromech.info

Web: www.electromech.info

Focuz Infotech

Focuz Infotech Advanced Education is the quality symbol of high-end Advanced Technology Education in the state. We are providing excellent services on Linux Technology Training, Certifications and live projects to students and corporates, since 2000.

Cochin

Tel: 0484-2335324

Email: enquiry@focuzinfotech.com

Web: www.focuzinfotech.com

G-TEC Computer Education

ISO 9001:2000 certified IT Company, International Testing Centre, Specialised in Multimedia & Animation, conduct MCP, MCSE 2000, MCDBA and MCSA certificates, CCNA, CCNP, the Only authorized centre by INTERNATIONAL AND EUROPEAN COMPUTER UNION to conduct ICDDL, Adobe Certifications, training on Web Designing, Tally, Spoken English. Conducts Corporate and institutional training. International certifications issued.

Bangalore

Tel: 080-43567000

Email: gtec.indiranagar@gmail.com

Gujarat Infotech Ltd

GIL is a IT company and 17 years of experience in computer training field. We have experience and certified faculty for the open Source courses like Redhat, Ubuntu, and PHP, MySQL

Ahmedabad

Tel: 079-27452276, Fax: 27414250

Email: info@gujaratinfotech.com

Web: www.gujaratinfotech.com

Lynus Academy Pvt Ltd

India's premier Linux and OSS training institute.

Chennai

Tel: 044-42171278, 9840880558

Email: contactus@lynusacademy.com

Web: www.lynusacademy.com

Linux Learning Centre Private Limited

Pioneers in training on Linux technologies.

Bangalore

Tel: 080-22428538, 26600839

Email: info@linuxlearningcentre.com

Web: www.linuxlearningcentre.com

Maze Net Solutions (P) Ltd

Maze Net Solution (P) Ltd, is a pioneer in providing solutions through on time, quality deliverables in the fields of BPO, Software and Networking, while providing outstanding training to aspiring IT Professionals and Call Center Executives. Backed by a team of professional workforce and global alliances, our prime objective is to offer the best blend of technologies in the spheres of Information Technology (IT) and Information Technology Enabled Services (ITES).

Chennai

Tel: 044-45582525

Email: info@mazenetsolution.com

Web: www.mazenetsolution.com

Netweb Technologies

Simplified and scalable storage solutions.

Bangalore

Tel: 080-41146565, 32719516

Email: info@netwebindia.com

Web: www.netwebindia.com

New Horizons India Ltd

New Horizons India Ltd, a joint venture of New Horizons Worldwide, Inc. (NASDAQ: NEWH) and the Shriram group, is an Indian company operational since 2002 with a global foot print engaged in the business of knowledge delivery through acquiring, creating, developing, managing, lending and licensing knowledge in the areas of IT, Applied Learning. Technology Services and Supplementary Education. The company has pan India presence with 15 offices and employs 750 people.

New Delhi

Tel: 011-43612400

Email: info@nhindia.com

Web: www.nhindia.com

Network NUTS

India's only Networking Institute by Corporate Trainers. Providing Corporate and Open classes for RHCE / RHCSA training and certification. Conducted 250+ Red Hat exams with 95% result in last 9 months. The BEST in APAC.

New Delhi

Tel: 46526980-2

Mobile: 09310024503, 09312411592

Email: info@networknuts.net

Web: www.networknuts.net

STG International Ltd

An IT Training and Solution Company. Over an experience of 14 years. We are ISO 9001:2000 Certified. Authorised Training Partners of Red Hat & IBM-CEIS. We cover all Software Trainings.

New Delhi

Tel: 011-40560941-42, Mobile:

09873108801

Email: rakhi@stg.in

Web: www.stgonline.com

www.stgglobal.com

FOSS Yellow Pages

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

TNS Institute of Information Technology Pvt Ltd

Join RedHat training and get 100% job guarantee. World's most respected Linux certification. After RedHat training, you are ready to join as a Linux Administrator or Network Engineer.

New Delhi

Tel: 011-3085100, Fax: 30851103

Email: nks@tiit.co.in

Web: www.tiit.co.in

Webel Informatics Ltd

Webel Informatics Ltd (WIL), a Government of West Bengal Undertaking. WIL is Red Hat Training Partner and CISCO Regional Networking Academy. WIL conducts RHCE, RHCSS, CCNA, Hardware and Software courses.

Kolkata

Tel: 033-22833568, Mobile: 09433111110

Email: enquiry@webelinformatics.com

Web: www.webelinformatics.com

Training for Professionals

AEM

AEM is the Best Certified Redhat Training Partner in Eastern India since last 3 years. AEM conducted more than 500 RHCE exams with 95-100% pass rate. Other courses—RHCSS, SCNA, MCSE, CCNA.

Kolkata

Tel: 033-25488736, Mobile: 09830075018

Email: sinhatuhin1@gmail.com

Web: www.aemk.org

Agam Institute of Technology

In Agam Institute of Technology, we provide hardware and networking training since last 10 years. We specialise in open source operating systems like Red Hat Linux since we are their preferred training partners.

Dehradun

Tel: 0135-2673712, Mobile: 09760099050

Web: www.agamtecindia.com

Amritha Institute of Computer Technology

Amritha Technologies provides an extensive training in high end certification programs and Networking Solutions like Redhat Linux, Redhat Security Services, Cisco, Sun Solaris, Cyber Security Program IBM AIX and so on with a strong focus on quality standards and proven technology processes with most profound principles of

Love and Selfless Service.

Mobile: 09393733174

Email: aict.hybd@amrita.ac.in

Web: www.amritahyd.org

Centre For Industrial Research and Staff Performance

A Unique Institute catering to the need for industries as well as Students for trainings on IT, CISCO certification, PLC, VLSI, ACAD, Pneumatics, Behavior Science and Handicraft.

Bhopal

Tel: 0755-2661412, 2661559

Fax: 0755-4220022

Email: crisp@crispindia.com

Web: www.crispindia.com

Center for Open Source Development And Research

Linux, open source & embedded system training institute and development. All trainings provided by experienced experts & administrators only. Quality training (corporate and individual). We expertise in open source solution. Our cost effective business ready solutions caters of all kind of industry verticals.

New Delhi

Mobile: 09312506496

Email: info@cfosdr.com

Web: www.cfosdr.com

Cisconet Infotech (P) Ltd

Authorised Red Hat Study cum Exam Centre. Courses Offered: RHCE, RHCSS, CCNA, MCSE

Kolkata

Tel: 033-25395508, Mobile: 09831705913

Email: info@cisconetinfn.com

Web: www.cisconetinfn.com

CMS Computer Institute

Red Hat Training partner with 3 Red Hat Certified Faculties, Cisco Certified (CCNP) Faculty, 3 Microsoft Certified Faculties having state Of The Art IT Infrastructure Flexible Batch Timings Available..Leading Networking Institute in Marathwada

Aurangabad

Tel: 0240-3299509, 6621775

Email: aurangabad@cmsinstitute.co.in

Web: www.cmsaurangabad.com

Cyber Max Technologies

OSS Solution Provider, Red Hat Training Partners, Oracle, Web, Thin Clients, Networking and Security Consultancy. Also available CCNA and Oracle Training on Linux. Also available Laptops & PCs

Bikaner

Tel: 0151-2202105, Mobile: 09928173269

Email: cmtech.bikaner@gmail.com,

kr.gupta.ashish@gmail.com

Disha Institute

A franchisee of Unisoft Technologies, Providing IT Training & Computer Hardware & Networking

Dehradun

Tel: 3208054, 09897168902

Email: thedishainstitute@gmail.com

Web: www.unisofttechnologies.com

EON Infotech Limited (TECHNOSchool)

TechnoSchool is the most happening Training Centre for Red Hat (Linux- Open Source) in the Northern Region. We are fully aware of the Industry's requirement as our Consultants are from Linux industry. We are committed to make you a total industry ready individual so that your dreams of a professional career are fulfilled.

Chandigarh

Tel: 0172-5067566-67, 2609849

Fax: 0172-2615465

Email: info@technoschool.net

Web: http://technoschool.net

GT Computer Hardware Engineering College (P) Ltd

Imparting training on Computer Hardware Networking, Mobile Phone Maintenance & International Certifications

Jaipur

Tel: 0141-3213378

Email: franchise_gt@gteducation.net

Web: www.gteducation.net

HCL Career Development Centre Bhopal

As the fountainhead of the most significant pursuit of human mind (IT), HCL strongly believes, "Only a Leader can transform you into a Leader". HCL CDC is a formalization of this experience and credo which has been perfected over three decades.

Bhopal

Tel: 0755-4094852

Email: bhopal@hclcdc.in

Web: www.hclcdc.in

IINSTRIX E Technologies Pvt Ltd

No. 1 Training prvider in this region.

meerut

Tel: 0121-4020111, 4020222

Mobile: 09927666664

Email: jai@iintrix.com

Web: www.iintrix.com

Indian Institute of Job Oriented Training Centre

Ahmedabad

Tel: 079-40072244—2255—2266

Mobile: 09898749595

Email: info@iijt.net

Web: www.iijt.net

Institute of Advance Network Technology (IANT)

•Hardware Engg. •Networking
•Software Engg. •Multimedia Training.

Ahmedabad

Tel: 079-32516577, 26607739

Fax: 079-26607739

Email: contact@iantindia.com

Web: www.iantindia.com

IPCC

Bridging Gap with professionals.

Lucknow

Tel: 0522-3919496

Email: ipcclo@yahoo.co.in

Web: www.ipcc.co.in

Koenig Solutions (P) Ltd

A reputed training provider in India. Authorised training partner of Red Hat, Novell and Linux Professional Institute. Offering training for RHCE, RHCSS, CLP, CLE, LPI - 1 & 2.

New Delhi

Mobile: 09910710143, Fax: 011-25886909

Email: info@koenig-solutions.com

Web: www.koenig-solutions.com

NACS/CIT

We are Providing Training of LINUX to Professional & Cooperate.

Meerut

Tel: 0121-2420587, Mobile: 9997526668

Email: Info@nacsglobal.com

Web: www.nacsglobal.com

NACS Infosystems (P) Ltd

NACS is a organization which is providing training for all international certification, and also NACS is the authorized Training Partner of Redhat and also having testing centre of THOMSON PROMETRIC and PEARSON VUE.

To advertise in this section,
please contact

Somaiah (B'lore:

09986075717) Dhiraj (Delhi:

09811206582) on

011-2681-0602 Extn. 222

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

Meerut

Tel: 0121-2767756, Fax: 0121-4006551
Mobile: 09897796603
Email: info@nacsnglobal.com,
mohit@nacsnglobal.com.
Web: www.nacsnglobal.com

Netdiox Computing Systems

We are one-of-a-kind center for excellence and finishing school focusing on ground breaking technology development around distributed systems, networks, storage networks, virtualisation and fundamental algorithms optimized for various appliance.

Bangalore

Tel: 080-26640708
Mobile: 09740846885
Email: info@netdiox.com

NetMax-Technologies

Training Partner of RedHat, Cisco

Chandigarh

Tel: 0172-2608351, 3916555
Email: mail.netmax@gmail.com
Web: www.netmaxtech.com



Netxpert institute of Advance Networking

Netxpert Noida is a Leading organization to provide Open Source training on RedHat Linux RHCT and RHCE Training with 30Hrs. extra exam preparation module.

Noida

Tel: 0120-4346847, Mobile: 09268829812
Email: anoop.vyas@netxpertindia.com
Web: www.netxpertindia.com

Netzone Infotech Services Pvt Ltd

Special batches for MCSE, CCNA and RHCE on RHEL 5 with exam prep module on fully equipped labs including IBM servers, 20+ routers and switches etc. Weekend batches are also available.

New Delhi

Tel: 011-46015674, Mobile: 9212114211
Email: info@netzoneindia.net

Neuron IT Solutions

We offer end to end services and support to implement and manage your IT Infrastructure needs. We also offer Consulting services and Training in Advanced Linux Administration.

Chennai

Mobile: 09790964948
Email: enquiry@neuronit.in
Web: www.neuronit.in

Plexus Software Security Systems Pvt Ltd

Plexus, incorporated in January 2003 is successfully emerged as one of the best IT Company for Networking, Messaging & Security Solutions and Security Training. Networking, Messaging & Security solutions is coupled with the expertise of its training; this has put Plexus in the unique position of deriving synergies between Networking, Messaging & Security Solutions and IT Training.

Chennai

Tel: 044-2433 7355
Email: training@plexus.co.in
Web: www.plexus.co.in

Professional Group of Education

RHCE & RHCSS Certifications

Jabalpur

Tel: 0761-4039376,
Mobile: 09425152831
Email: naidu.vikas@gmail.com

Q-SOFT Systems & Solutions Pvt Ltd

Q-SOFT is in a unique position for providing technical training required to become a Linux Administration under one roof. Since inception, the commitment of Q-SOFT towards training is outstanding. We Train on Sun Solaris, Suse Linux & Redhat Linux.

Bangalore

Tel: 080-26639207, 26544135, 22440507
Mobile: +91 9945 282834

Email: counsellors@qsoftindia.com
Web: www.qsoftindia.com

Software Technology Network

STN is one of the most acknowledged name in Software Development and Training. Apart from providing Software Solutions to various companies, STN is also involved in imparting High-end project based training to students of MCA and B.Tech etc. of various institutes.

Chandigarh

Tel: 0172-5086829
Email: stn2001@rediffmail.com
Web: stntechnologies.com

South Delhi Computer Centre

SDCC is for providing technical training courses (software, hardware, networking, graphics) with career courses like DOEACC "O" and "A" Level and B.Sc(IT), M.Sc(IT), M.Tech(IT) from KARNATAKA STATE OPEN UNIVERSITY.

New Delhi

Tel: 011-26183327, Fax: 011-26143642
Email: southdelhicomputercentre@gmail.com, southdelhicomputercentre@hotmail.com.
Web: www.itwhizkid.com
www.itwhizkid.org

Ssystems Quest

Making Tomorrow's professionals TODAY

Bangalore

Tel: 080-41301814
Email: directorv@ssystemsquest.com
Web: www.ssystemsquest.com

Trimax FuturePerfect

A Div of Trimax IT Infrastructure and Services Limited. Redhat RHCE, RHCT Training & Exam Center, MCTS, MCITP, MCSE 03, CCNA, CCNP, Prometric Center.

Mumbai

Tel: 022-40681313, Mobile: 09987705638
Fax: 022-40681001

Email: futureperfect@trimax.in
Web: www.trimax.in

Vibrant e Technologies Ltd

Vibrant e Technologies Ltd. Is a authorised Red Hat Test and Testing Centre, has won the prestigious award "REDHAT BEST CERTIFIED TRAINING PARTNER 2007-2008" for Western region. Vibrant offers courses for RHCE 5, RHCSS etc.

Mumbai

Tel: 022-26285066/6701
Email: vibrant@vsnl.net
Web: www.vibrantcomputers.com

Ultramax Infonet Technologies Pvt Ltd

Training in IT related courses and authorised testing center of Prometric, Vue and Red Hat.

Mumbai

Tel: 022-67669217
Email: unmesh.raote@ultramaxit.com
Web: www.ultramaxit.com

Yash Infotech

Authorized Training & Exam Center. Best Performing Center in Lucknow for RH Training and Examinations. LINUX & Open Source training institute for IT professionals & Corporate Offering Quality Training for RHCE, RHCSS, PHP, Shell Script, Virtualization and Troubleshooting Techniques & Tools.

Lucknow

Tel: 0522-4043386, Fax: 0522-4043386
Email: yashinfotech.lko@gmail.com

Web Hosting

Perfect Innovation

Web Hosting Spider

- Web Hosting • Web Design
- Web Application Development
- SMS Hosting • Corporate Hosting
- Dedicated Servers

Puducherry

Tel: 0413-3202726, 3246999
Email: admin@webhostingspider.com
Web: www.webhostingspider.com

Want to register your organisation in
FOSS Yellow Pages For **FREE***

Call: Dhiraj (Delhi) 09811206582 Somaiah (Bangalore) 09986075717
or mail: dhiraj.khare@efyindia.com, somaiah.km@efyindia.com

*Offer for limited period.

